BrightStor® CA-ASTEX™ Performance

Administrator Guide

2.8

Computer Associates™

Contents

Chapter 1: Introduction

Chapter 2: PMC: Planning Before Customizing

Chapter 3: MMC: Planning Before Customizing

Chapter 4: Configuring and Starting CA-ASTEX

Chapter 5: Tailoring CA-ASTEX Parameters

Appendix A: PMC: Calculating Storage and Space Requirements

Appendix B: MMC: Calculating Storage and Space Requirements

Appendix C: PMC Technical Descriptions

Appendix D: Contacting Technical Support

Index

Chapter 1

Introduction

This guide describes how to configure and customize BrightStor CA-ASTEX Performance (CA-ASTEX) to run at your site. It covers these topics:

- Configuring CA-ASTEX
- Planning before customizing
- Tailoring CA-ASTEX parameters
- Starting CA-ASTEX
- Calculating storage and space requirements
- Troubleshooting
- Technical descriptions

CA-ASTEX Documentation

The following manuals are included with the CA-ASTEX product:

Getting Started

Administrator Guide

User Guide

Storage Manager User Guide

Cache Manager User Guide

DASD Manager User Guide

Migration Manager User Guide

Message Guide

Contacting Technical Support

For technical assistance with this product, please contact Computer Associates Technical Support at http://esupport.ca.com. Technical support is available 24 hours a day, 7 days a week.

Before contacting technical support, please try to resolve your problem by reading this manual and related publications, and by using any Help systems associated with the product.

Chapter

2

PMC: Planning Before Customizing

Identifying the Volumes to be Monitored

By default, CA-ASTEX monitors performance for *all* of your online OS/390 volumes. In addition, CA-ASTEX dynamically detects when a volume is varied online or offline, renamed, or when the defined SMF storage group changes, and begins monitoring the volume's performance.

However, if you want it to monitor just *selected* volumes, you can list those volumes in the DEFINE VOLUMES section of your parameter data set (the ASXPARM member of ASX.CNTL).

You can also have CA-ASTEX monitor performance for certain groups of volumes:

Contention Groups

Volumes that share a common I/O path form a *contention group*, because their I/O requests "contend" for the same I/O channels, controllers, and strings.

By default, CA-ASTEX assigns volumes to contention groups in both single- and multiple-CPU environments. The contention group names are based on the Storage Director IDs in the Storage Director of the associated DASD controller. In multiple-CPU environments, the contention groups are assigned across all CPUs.

You should rely on CA-ASTEX's default contention group assignments rather than creating your own. However, you can specify parameters for these contention groups in the DEFINE CGROUP section of the parameter data set.

Storage Groups

Volumes that contain similar data or provide service to the same applications can be members of a *storage group*. One or more storage groups must be defined in your system if you want to use CA-ASTEX's Storage Performance Expert, because it recommends data set moves only between volumes in the same storage group.

- If you are using SMS, CA-ASTEX automatically assigns volumes to the correct SMS storage groups. You should rely on these default storage group assignments.
- If you are not using SMS, you must define your own storage groups in the DEFINE SGROUP section of your parameter data set.

Logical Volume Groups

The set of logical volumes that map or are defined to the same physical disk(s) make up a ***Logical Volume Group*** (LVG). I/O to these logical volumes are contending for the same physical disk(s). Delays occur at the logical volume group level when multiple logical volumes in a LVG are too busy, when data on one or more logical volumes is not cache-friendly, or when the LVG's volumes have a high write-to-read I/O ratio.

CA-ASTEX automatically determines what volumes belong in a logical volume group (LVG) and names that group. For IBM's RAMAC DASD, a logical volume group represents a drawer. For EMC Symmetrix ICDAS, a logical volume group represents the set of hyper-volumes defined to a physical disk.

Defining an I/O Response Time Objective

Since CA-ASTEX monitors performance by measuring I/O response times, you need to define *I/O response time objectives* that CA-ASTEX can use to distinguish between good and bad performance. You can set response time objectives at the system, contention group, storage group, and volume levels through parameters in the DEFINE PARMS, DEFINE CGROUP, and DEFINE SGROUP sections of your parameter data set.

How CA-ASTEX Sets Response Time Objectives

CA-ASTEX provides default millisecond response (MSR) time objectives for SMS and non-SMS environments. SMS data sets are defined to three basic storage classes: must-cache, may-cache, or never-cache. CA-ASTEX specifies default MSR objectives for these classes:

- 10 ms 90% of the time for must-cache data sets
- 25 ms 90% of the time for may-cache data sets
- 35 ms 90% of the time for never-cache data sets

For non-SMS data sets, CA-ASTEX specifies an MSR objective of 30 ms 90% of the time.

Specifying Your Own Response Time Objectives

If you want to specify your own MSR objectives for your SMS or non-SMS environments, consider the following:

- Set response time objectives according to workload priorities. The relative difference among response time objectives in relation to these priorities is more important than the actual values of the objectives. For example, if IMS is the most critical workload you have, its response time objective should be lower than all other response time objectives.
- Start with conservative objectives. Initially, you might want to specify higher response time objectives (for example, 50 milliseconds instead of 30 milliseconds). That way, only the very worst performance problems show up, and you will not have to face an overabundance of data at the start.
- If you are not sure what "good" response time objectives might be, run CA-ASTEX with the defaults. Observe the average I/O response times, and base your response time objectives on those averages.
- If you want to specify SMS and non-SMS MSR objectives, you should:
  - Specify SMSPARMS=YES

  - Use the MUSTMSR, MAYMSR, and NEVMSR parameters to specify response objectives for SMS volumes
  - Use the NSMSMSR parameter to specify response objectives for non-SMS volumes
- If you want to disregard the SMS MSR objective values and define one response time objective, you should:
  - Specify SMSPARMS=NO in the default parameters section
  - Use the NSMSMSR parameter to specify response objectives for your volumes

Regulating the Collection of Detailed Data

To regulate the amount of detailed job and data set performance data that CA-ASTEX collects and maintains, you define an appropriate *data collection mode*:

Detail Mode	CA-ASTEX continuously gathers detailed job and data set data for a particular volume or group of volumes.
Exception Mode	CA-ASTEX collects detailed data only when I/O response time exceeds the defined response time objective. This is the default.

You can specify data collection modes at the system, contention group, storage group, and volume levels by using one or more IMODE parameters in the DEFINE PARMS, DEFINE CGROUP, and DEFINE SGROUP sections of your parameter data set. You can also dynamically change the data collection mode for individual volumes and groups of volumes with the SVOL and SGRP operator commands. See the chapter "Understanding the Migration Manager" in the CA-ASTEX *User Guide* for more information.

Consider the following before specifying or changing a data collection mode:

- **The mode controls data quantity and "mix" of detailed data**.

  When running in exception mode, CA-ASTEX collects and saves only job and data set data that is associated with periods of poor performance. On the other hand, in detail mode, CA-ASTEX collects data for periods of both good and bad performance. When you run CA-ASTEX in detail mode, you could find yourself with more data than you need.

- **The modes should be consistent for volumes within and across groups**.

  Performance data can be misleading when CA-ASTEX collects some of it in exception mode and some in detail mode. So, you should be careful about specifying exception mode for some volumes in a group and detail mode for the others. Likewise, you should watch for the situation where the same volume belongs to both a contention group (where, for example, detail mode is defined) and a storage group (where exception mode is defined).

  Note: The mode assigned to a volume in a storage group overrides the mode assigned to the same volume in a contention group.

- **The mode affects CA-ASTEX's solutions to your problems**.

  CA-ASTEX may offer differing solutions when collecting data in exception mode as opposed to detail mode. For instance, the Volume Reorganization utility might position a volume's data sets differently if it considers data collected only during periods of poor performance instead of data collected continuously. In the same situation, the solution plan might suggest different moves to optimize or balance path and volume performance.

Regulating CPU Usage

By design, CA-ASTEX can run continuously while using only a very small percentage of the CPU. However, CA-ASTEX's CPU usage increases proportionally as it collects more detailed data. To lower the amount of overhead CA-ASTEX incurs, you can:

- Change the data collection mode for some of your volumes from detail mode (IMODE=DM) to exception mode (IMODE=EM).
- When running in exception mode, set conservative response objectives with the NSMSMSR, MUSTMSR, MAYMSR, and NEVMSR parameters.
- Decrease the number of I/O requests being monitored. For example, instead of having CA-ASTEX monitor every I/O request (IOSAMPL=1), you could have it monitor every other I/O request (IOSAMPL=2). Changing the IOSAMPL parameter should be done as a last resort. If you change the value of the IOSAMPL parameter to any value other than 1, be sure to turn the Cache Optimizer off.

Deciding How to Store CA-ASTEX Data

You must decide whether or not to have CA-ASTEX write performance data. You can use both the Interval Database and the SMF data set, but this involves additional space and overhead requirements. We recommend one of the following:

- **A CA-ASTEX Interval Database**

  This option gives you immediate online access to detailed performance data and recommended solutions to problems. You can also use this database as input to batch reports or the merge utility, so you can produce reports covering your entire complex and view recommended solutions online.

- **Your SMF data set**

  This option requires that you convert the SMF data set to an Interval Database (using the CA-ASTEX merge utility) to gain access to the CA-ASTEX online screens and expert analysis.

Calculating Your Common Storage Needs

The amount of common storage space CA-ASTEX requires depends upon the number of volumes it has to monitor. For most of its storage, CA-ASTEX uses extended common storage (above 16 megabytes in virtual storage) if it is available. You can use the CLIST in the ASXCALC member of the data set to determine your approximate common storage needs, or you can calculate them yourself by following the instructions in the appendix "PMC: Calculating Storage and Space Requirements."

Chapter 3

MMC: Planning Before Customizing

Importance of the PREFIX Parameter

The PREFIX parameter is described in the chapter "Tailoring CA-ASTEX Parameters," with the other ASXPARM keywords. Because it requires your attention at the planning stage, PREFIX is introduced here. The PREFIX=*prefix* parameter establishes a high-level qualifier for all the permanent data sets, including the database, used by the Migration Manager through dynamic allocation.

Ensure MMC Can Obtain Sufficient Temporary Space

The Migration Manager is an application product that can require very large quantities of space for intermediate data sets during its daily cycle. MMC design goals are to maximize the use of dynamic allocation to shield you from having to manage JCL that could only estimate space requirements. MMC makes only one arbitrary guess (during the initial space inventory) about how much space is needed. All other MMC space requests are based on exact calculations for how much is really needed. However, a DASD complex with 2000 volumes and 450,000 data sets requires more intermediate space than a configuration half that size. The characteristics of the data sets are as follows:

- Data set names are *permanent* and formed as DSN=*prefix.ddname*, where:

  prefix is the value of the PREFIX parameter you choose

  ddname is the file identifier used by the MMC program

- Data set usage is *transient* in that data sets are deleted by MMC as soon as they are no longer needed by a phase of the daily cycle.
- The purpose of the permanent names for transient usage is to enable restart at a phase level if needed (by means of the MODIFY command and RESTART=STEP).

- Allocations use the RLSE parameter, so used space is released at CLOSE time.

  Note: An important exception is dynamic allocation of the SORTWK*xx* data sets discussed in more detail below.

- Allocations use UNIT=(*unitname*,*unitcount*), where *unitname* and *unitcount* come from the DYNUNIT parameter, described in the chapter "Tailoring CA-ASTEX Parameters."

Verify Requirements with Storage Administrator's ACS Routines

MMC cannot be denied this space or have its requests trimmed back by ACS routines (DATACLAS, STORCLAS, MGMTCLAS, storage group, and so forth.) or the equivalent. If that happens, operational problems occur. When you choose a value for the PREFIX parameter (see "Importance of the PREFIX Parameter" in this chapter) for a description, make sure you have cleared this with the storage administrators who maintain ACS routines (or the equivalent allocation control software). It is critical that the space is not released during the CA-ASTEX multiple opening and closing of these data sets, and that the blocksize is not changed from 4080 for certain data sets.

You must choose a value for PREFIX so that when allocation requests are made, they are not restricted.

If problems persist, you can have your ACS routine direct the allocation to a non-SMS volume pool; otherwise, there is no requirement for non-SMS residency for MMC data sets.

Example 1: DATACLAS Requirements and SORTWKxx Data Sets

To be a good user of available SORT workspace, MMC makes its allocation requests dynamically at the beginning of the daily cycle. In this chapter you learn how to control these amounts using parameters in ASXPARM.

The dynamic allocation requests are made only once and without RLSE specified, because the data sets are reused by MMC many times during the daily cycle. If RLSE were specified on the allocation request, then the size would be trimmed back at CLOSE time to whatever was the requirement of the first usage of SORT. This could later become a problem with other intermediate files that required larger SORT extents.

When MMC completes the daily cycle, all SORTWK*xx* data sets are deleted. The space is freed and is not required until the next START time for the daily cycle.

Limit PREFIX to a Maximum of 27 Characters

This should allow adequate description in the name, while permitting you to follow the recommendation of running the ASXMMGDG job, described in the chapter "Configuring and Starting CA-ASTEX." Even though the PREFIX value can be up to 36 characters, we recommend not exceeding 27 characters.

Synchronize PREFIX with the Value in the ALLOCMMC Job

This means that your tailoring of the JCL for the ALLOCMMC job (see the sections "Step 3: Decide Whether to Use the Quick-Start Facility for MMC" and "Using the Quick-Start Facility (MMC Only)," "Step 5: Allocate the New MMC Database" in the "Configuring and Starting CA-ASTEX" chapter) must specify the same PREFIX value as discussed above. The database is composed of several logical parts, and each data set name is of the form *prefix.ddname*. Unlike the intermediate (work) data sets that get deleted when no longer needed, the database is permanent and persists from day to day.

prefix.DB*xxxx* data sets do not need a secondary allocation, because they are pre-formatted by a separate utility (AMMFMTDB). During the daily cycle, a wrap occurs rather than an out-of-space condition, and there is never expansion to new extents. Allocation on different volumes would improve performance while complicating backup and recovery requirements. None of the database data sets should be multi-volume.

Likewise, when you tailor the JCL in AMMFMTDB to perform you use the same PREFIX value. Running this job causes the allocation to be always 100% used. For more information, see the sections "Step 3: Decide Whether to Use the Quick-Start Facility for MMC" and "Step 6: Format the New MMC Database" in the "Configuring and Starting CA-ASTEX" chapter.

Note: Make sure you complete the AMMFMTDB job shortly after ALLOCMMC, or you are subject to possible PARTIAL RELEASE of the database.

Exceptions and STOPX37 Considerations

prefix.FSRLOG begins with a very small allocation, expands as needed during the course of a day, and has automatic recovery from Sx37 abends. It contains the HSM SMF records (FSR) and eliminates the need for MMC to read dumped SYS1.MAN_ files from every OS/390 image with HSM. It is *never* deleted. Rather, when it has been successfully processed by the daily cycle, an EOF record is written at the beginning of the file, and the free space in this sequential data set is *not* released.

Special Note to Users of STOPX37: If you use the STOPX37 product, you should add Migration Manager to that product's exclude tables. This should be done by data set name (for example, *prefix*.FSRLOG), since this is the only MMC data set (other than the database) with its own recovery from Sx37 abends. This exclusion is necessary because STOPX37 intercepts the out-of-space abend before the CA-ASTEX recovery is allowed a chance to handle the condition. When control is finally passed to the CA-ASTEX recovery code designed specifically to expand the data set gracefully, the environment has been changed by STOPX37 such that normal recovery cannot be completed. Use the exclusion facilities in STOPX37 so that Migration Manager can process without intervention. (This cannot be done with insertion of a special DD name in the JCL because MMC has no JCL for its dynamically allocated data sets.)

prefix.VCCDATA holds the inventory of primary DASD and is first allocated with DISP=SHR. Thus, it is possible for you to preallocate this data set, but we do not recommend doing this. Normally, it does *not* exist, and DISP=(NEW) allocation begins with a (guess) request for 100 cylinders primary, 20 cylinders secondary. UNIT=(SYSALLDA,1) is used, unless you code the DYNUNIT parameter. If this request fails, the amount is halved to (50,10). RLSE is used to free space at CLOSE. CA-ASTEX uses half-track blocking on whatever device type is provided by OS/390 allocation.

Unless you specify MICS=YES, this data set is deleted as soon as it is processed by the next phase (that is, if you DO pre-allocation, you must specify **MICS=YES** to keep the data set from being deleted, or arrange for the pre-allocation to be repeated daily). For example, you could arrange for an allocation type job to be submitted just before the START time on a daily basis. Yet these arrangements are exactly the type of concerns CA-ASTEX intended to relieve with dynamic allocation.

prefix.VCCHSM holds the inventory of ML1 and part of ML2, and is the second and last data set whose size we guess. All specifics are the same as for *prefix*.VCCDATA. Allocation and subsequent deletion for VCCHSM occurs concurrently with that of VCCDATA, so this is a short-term *peak* requirement for space.

Dynamically Allocated Data Sets for MMC

The following is a list of the data sets dynamically allocated:

***prefix*.DBMAIN**	Migration Manager Database (storage group and volume information)
***prefix*.DBFILE**	Migration Manager Database (data set information)
***prefix*.DBHIST**	Migration Manager Database (name and age history tables)
***prefix*.DBMISC**	Migration Manager Database (miscellaneous objects)
***prefix*.DBMOVE**	Migration Manager Database (move request queue)
***prefix*.FSRLOG**	Accumulated FSR activity data
***prefix*.VCCDATA**	Inventory of VTOC, VVDS, and SMS data
***prefix*.VCCHSM**	Inventory of HSM ML1 and ML2 migrate data
***prefix*.VCCNTROL**	Control file with object counts and time stamps

You should choose a prefix value that enables these data sets to be allocated on DASD that is common to all systems in a shared configuration. For example, do *not* use **PREFIX=SYS1.ASTEX**, because that would require cataloging in a system-specific catalog.

GDGs Created for MMC

Optionally, if you chose to tailor and submit the ASXMMGDG job as described in "Step 2: Perform Setup Tasks for MMC" in the "Configuring and Starting CA-ASTEX" chapter of this guide, a series of GDGs are created:

***prefix*.VCCHSMSG**	Summary of HSM Inventory Task
***prefix*.VCCSTATS**	Summary of Storage Inventory Results
***prefix*.VCCVTMSG**	Log of VTOC Scanning Subtask Activity
***prefix*.VCCVSMSG**	Log of VVDS Scanning Subtask Activity

Security Package Considerations

Normally, most sites establish a default user ID for started tasks with no privileges to prevent unauthorized users from having PROCs with inappropriate performance or security access. Therefore, you must ensure that when the START command is entered for the PROCMMC PROC, it is established with a user ID with these privileges:

- READ access to all VTOCs and SYS1.VVDS.* data sets
- ALTER authority for the MASTER catalog

  Note: Be assured that the MMC application does not change anything in a catalog and that all access is in fact, read-only. However, the Master Catalog is located off of system control blocks, dynamically allocated, OPENed and read sequentially as a KSDS. This process requires setting the "STEPCAT" flag in the PCCB while in protect key 0, which in turn requires ALTER authority to the MASTER catalog.
- READ access to the HSM PARMLIB data set
- READ access to the HSM migration control data set (MCDS)
- ALTER authority for all prefix.ddname data sets used by MMC

Allow Universal READ Access to prefix.DB* Data Sets

In establishing special privileges listed above for the application PROCMMC, make sure you allow READ access to the MMC database for people who want to use the ISPF dialog and for batch jobs that may need to access the database.

Enabling HSM to Log SMF Data

Migration Manager requires SMF data generated by HSM. This data is called *functional statistics records* (FSRs), and you may need to change your HSM parameters if you are not already generating the data.

Check the current ARCCMD*xx* member of your parameter data set used by HSM. Scan for the keyword **SMF** and ensure that it specifies **SMF(*nnn*)** rather than **NOSMF**. The Migration Manager parses this member during initialization and captures the number, but you get an error message and immediate termination if this is not set properly.

If you prefer to discard the SMF data generated by HSM, you can do so by means of the KEEPSMF=NO parameter discussed in the chapter, "Tailoring CA-ASTEX Parameters." The Migration Manager only needs HSM to generate the data, and you can tell the Migration Manager how to dispose of the record(s). We recommend that you let the data be logged to SMF (that is, specify **KEEPSMF=YES**).

Important Information About LNKLSTxx

CA-ASTEX installation requires that all of MMC and much of PMC be executed from an APF-authorized library. If you choose to make the APF-authorized library part of the LNKLST*xx* concatenation, you need to be aware of possible problems with the NeuMICS DASD Space Analyzer product (VCC).

Note: If the CA-ASTEX library is *not* in the LNKLST, or if you do *not* use the NeuMICS DASD product, then you can ignore this section.

VCC modules are a "common component" between CA-ASTEX and NeuMICS. Because the maintenance schedules for each product is different, it is best if each "user" of the VCC module set (that is, CA-ASTEX and NeuMICS) maintain separate APF libraries corresponding to the level of the product. This is much easier to maintain if each product is established using a //STEPLIB DD pointing to a library named in the IEAAPF*xx* list. This eliminates any concern over which product most recently installed replacement modules into a library named in the LNKLST*xx* list.

Note: NeuMICS clients can optionally have the daily inventory of DASD be triggered by the CA-ASTEX Migration Manager at the specified START time by coding **MICS=YES** in the ASXPARM data set and following other guidelines in the VCC6090 product change text.

MMC and CA-Multi-image Manager

If you do not have CA-Multi-image Manager, or if MMC does *not* run in a multi-image environment, then you can ignore this section.

Check with the person responsible for CA-Multi-image Manager and determine the "mode" of processing for the CA-Multi-image Manager Global Data Integrity Facility (GDIF). The mode of processing for GDIF is specified via the PROCESS= parameter on the GDIINIT statement found in the CA-Multi-image Manager initialization parameter member.

If **PROCESS=ALLSYSTEMS** is specified, then no change is required.

If **PROCESS=SELECT** is specified, then the following steps are required:

1. Review the CA-Multi-image Manager documentation regarding creating a qname list.
2. Add the following to the list of qnames:

```
ASTEXMM GDIF=YES SCOPE=SYSTEMS EXEMPT=NO ECMF=NO TRACE=NO RESERVES=KEEP
```

After version 2.8 is installed on *all* systems on which the Migration Manager runs, you can use **RESERVE=CONVERT** instead of **RESERVE=KEEP** to eliminate the hardware reserve done for the device on which *prefix*.DBMOVE exists.

Chapter 4

Configuring and Starting CA-ASTEX

Licensing Requirements

CAIRIM

CAIRIM, the Resource Initialization Manager, is used to provide product licensing for CA-ASTEX. CAIRIM is one of the common services of Unicenter Common Services for OS/390 (formerly Unicenter TNG Framework for OS/390). CAIRIM prepares your operating system environment for Computer Associates OS/390 products and components and executes them.

You must use either the Unicenter Common Services for OS/390 (any version) or CA90s Services for MVS version 1.0 level 9611 or higher.

CAIRIM routines are grouped under CA MVS Dynamic Service Code S910. Review the Unicenter Common Services for OS/390 documentation for further details about the features and associated utilities of CAIRIM.

CA LMP

CA-ASTEX requires CA LMP (License Management Program) to initialize correctly. CA LMP also provides a standardized and automated approach to the tracking of licensed software. CA LMP is provided as an integral part of CAIRIM.

LMP Key Certificate

Examine the CA LMP Key Certificate you received with your CA-ASTEX installation or maintenance tape. Your certificate contains the following information:

Term	Definition
Product Name	The trademarked or registered name of CA-ASTEX as licensed for the designated site and CPUs.
Product Code	A two-character code that corresponds to the CA-ASTEX product.

Term	Definition
Supplement	The reference number of your license for the particular CA-ASTEX facility in the format: `nnnnnn - nnn.` This format differs slightly inside and outside North America, and in some cases may not be provided at all.
CPU ID	The code that identifies the specific CPU for which installation of CA-ASTEX is valid.
Execution Key	An encrypted code required by CA LMP for CA-ASTEX installation. During installation, it is referred to as the "LMP Code."
Expiration Date	The date your license for CA-ASTEX expires. It is of the format: ddmmmyy, as in 15JUN01.
Technical Contact	The name of the designated technical contact at your site responsible for installation and maintenance of CA-ASTEX. This is the person to whom CA addresses all CA LMP correspondence.
MIS Director	The name of the Director of MIS or the person who performs such a function at your site. If only the title is indicated on the certificate, you should supply the responsible individual's name when correcting and verifying the certificate.
CPU Location	The address of the building in which the CPU is located.

Specifying the LMP Code

You must add the CA LMP Execution Key provided on the Key Certificate to the CAIRIM parameters to ensure proper initialization of CA-ASTEX. To define a CA LMP Execution Key to the CAIRIM parameters, modify member KEYS in CAI.PPOPTION.

The statement structure for member KEYS is:

```
PROD(pp) DATE(ddmmmyy) CPU(tttt-mmmm/ssssss) LMPCODE(kkkkkkkkkkkkkkkk)
```

Parameter definitions are as follows:

Parameter	Definition
pp	One or more two-character Product Codes. This code agrees with the Product Code already in use by the CAIRIM initialization parameters for any earlier versions of CA-ASTEX (if applicable). This is required. The following is a list of the CA-ASTEX LMP codes: **D0** CA-ASTEX Complete (four components) **EQ** CA-ASTEX DASD Manager **D7** CA-ASTEX CACHE Manager **ER** CA-ASTEX Storage Manager **FT** CA-ASTEX Migration Manager
ddmmmyy	The CA LMP licensing agreement Expiration Date.
tttt-mmmm	The CPU Type and Model on which CA LMP is to run (for example, 3090-600). If the CPU Type and/or Model require less than four characters, blank spaces are inserted for the unused characters. This is required.
ssssss	The Serial Number of the CPU on which CA LMP is to run. This is required.
kkkkkkkkkkkkkkk	The Execution Key needed to run CA LMP. This CA LMP Execution Key is provided on the Key Certificate shipped with each CA LMP software solution.

The following is an example of a control statement for the CA LMP execution software parameter. The Product Code and Execution Key value is different when you install CA-ASTEX at your site.

```
PROD(D0) DATE(27JAN02) CPU(3090-600 /370623) LMPCODE(52H2K06130Z7RZD6)
```

For a full description of the procedure for defining the CA LMP Execution Key to the CAIRIM parameters, refer to the Unicenter Common Services for OS/390 documentation.

Performing Configuration Steps

The following table summarizes the procedure you follow to configure CA-ASTEX. These steps are described in detail in subsequent sections of this chapter.

Step	Action
1	Allocate PMC data files.
2	Perform setup tasks for MMC.
3	Decide whether to use the Quick-Start facility for MMC.
4	Tailor the ISPF options.
5	Tailor and submit the ASXFMTDB job to initialize the PMC databases.
6 (Optional)	Install the ARCRDEXT exit to optimize non-directed DFHSM recalls of data sets that are not managed by SMS. (This exit is for PMC only.)
7 (Optional)	Set up connectivity to BrightStor DASD Manager Option for Unicenter and CA-ASTEX (formerly Unicenter TNG DASD Manager Option).
8 (Optional)	Set up security for the Command Interface.
9	Customize the PMC start-up procedure.
10	Customize the MMC start-up procedure.

Step 1: Allocate PMC Data Files

The ALLOCPMC job found in ASX.CNTL allocates data sets for CA-ASTEX to use, as follows:

- A data set(s) that serves as the Interval Database
- A data set that serves as the Trend Database
- A data set that serves as the Configuration File
- An optional user-defined data set for CA-ASTEX trace data

The following pages describe some of the special considerations you must make when you allocate these data sets.

Allocating an Interval Database

An Interval Database stores data collected by the CA-ASTEX PMC. Using an Interval Database, instead of your SMF data set, gives you immediate on-line access to this data.

You can allocate your Interval Databases in two ways:

1. **One Interval Database that consists of one file for each system on which CA-ASTEX is running.** Volume, data set, and job data are written to this one file. *Do not* give this file a low-level qualifier of .SUM. (See information below for an explanation of how to use the .SUM qualifier.)
2. **One Interval Database that consists of two files for each system on which CA-ASTEX is running.** One file contains your volume (summary) data and the other contains data set and job (detail) data. Because detail data usually consists of large numbers of records per interval, keeping it in one file with your summary data usually causes the file to wrap more frequently.

   By separating the two types of data, you can accumulate more intervals of summary data for performance analysis.

   When you name the summary file, give it a low-level qualifier of .SUM. The detail file must have a low-level qualifier of .DET. Both file names must have the same high-level qualifier.

   Note: If you delete and reallocate either the .DET file or the .SUM file, you must delete and reallocate the other one as well. CA-ASTEX treats the .DET and .SUM files as one Interval Database.

You can use the CLIST in the ASXCALC member of ASX.CNTL to determine your space requirements for an Interval Database, or you can calculate them yourself by following the instructions in the Appendix "PMC: Calculating Storage and Space Requirements."

Important: *The maximum size of the Interval Database should not exceed 2910 cylinders.*

Allocating a Trend Database

The Trend Database stores long-term performance trends accumulated from interval data in the Interval Database or the SMF data set. The Storage Performance Expert uses this trend data to produce a complete solution plan. The Cache Management Expert uses trend data to solve cache utilization problems that are occurring across an entire production cycle. The Allocation Director uses it as input to make its allocation decisions. One Trend Database should be used for all systems that share the same DASD.

You can use the CLIST in the ASXCALC member of ASX.CNTL to determine your space requirements for a Trend Database, or you can calculate them yourself by following the instructions in the Appendix, "PMC: Calculating Storage and Space Requirements."

Allocating a Data Set for the Configuration File

If you plan to use the CA-ASTEX Excluded Data Set List or any of the Configuration Manager functions available with CA-ASTEX, you need to allocate a Configuration File. This file contains:

- Information about all configuration changes developed by the Configuration Manager
- A list of data sets and volumes to be excluded from the Storage Performance and Cache Management Experts' analyses
- A list of data sets and volumes to be included in or excluded from Cache Simulation
- A list of data sets to be re-blocked and/or excluded from the Device Migration Utility

The Configuration File should have a physical sequential organization. Allocate at least four tracks of DASD space. One Configuration File may be needed for each user of CA-ASTEX.

Allocating a Data Set for the Trace File

You must have CA-ASTEX write its trace records to a user-defined data set rather than a SYSOUT data set if you expect to perform any of these tasks:

- Use trace data as input to the Workload Reporter
- Use trace data for Cache Simulation
- Use trace data for the TTR Analysis Report or PDS reorganization
- Run trace data through your own trace analysis program

To allocate a user-defined data set for trace records, calculate your space requirements, taking into account the following:

- CA-ASTEX creates a record for the trace data set every time it traces an I/O request.
- The trace record is a variable length record that can be a maximum of 302 bytes long. The length varies depending on the data compression techniques used by CA-ASTEX.
- You should allocate 10 to 100 cylinders of space, depending on how much I/O tracing you anticipate CA-ASTEX to do.

Step 2: Perform Setup Tasks for MMC

This step is necessary only if you are running MMC.

At this point, various parameters need to be selected and a PROC needs to be tailored. To proceed, you should take time to review the issues discussed in the chapter "MMC: Planning Before Customizing."

To proceed with the utilities, you need to have tailored the following:

- PROCMMC PROC

  See "Step 10: Tailor the MMC Start-up Procedure" in this chapter for details on tailoring this cataloged procedure.

- ASXPARM data set

  All MMC parameters are described in shaded boxes and are listed in alphabetical order along with the PMC-only parameters. Some of the configuration related parameters are common to both components. See the chapter "Tailoring CA-ASTEX Parameters" for details on the meaning and syntax of the parameters in this data set.

- ASXMMGDG job

  To make troubleshooting easier, you should tailor and submit the ASXMMGDG job. This job allocates generation data groups (GDGs) for Migration Manager data sets containing VTOC and VVDS data. These generation data groups are easier to access and review than a SYSOUT data set.

  Make sure that the PREFIX parameter you specify for MMC matches the value you specify on the IDCAMS control statements. See the chapter "MMC: Planning Before Customizing" for additional information.

Step 3: Decide Whether to Use the Quick-Start Facility for MMC

This step is necessary only if you are running MMC.

Four categories of users are described in this section. The issue is whether the Quick-Start facility is to be used. If it is not, certain steps described in the Quick-Start section are repeated here for your convenience. We recommend that you read the introductory part of the Quick-Start section before you decide whether to perform Quick-Start processing. See the section, "Using the Quick-Start Facility (MMC Only)" in this chapter for more information.

Case 1: New or Existing User of MMC (Perform Quick-Start)

To the New User, we recommend that you perform Quick-Start processing. If you have been running MMC for less than six months, we recommend that you consider performing Quick-Start processing.

Case 2: New or Existing User of MMC (Do Not Perform Quick-Start)

There are a few reasons why you might want to skip Quick-Start:

- You have not been logging HSM FSR activity records to SMF.
- You have changed the SMF record number for the FSR records on a given SYSID in the period of time being considered.
- You have been running MMC for more than six months.
- You have experienced one or more of the following conditions within the last six months:
  - Large increases or decreases in DASD capacity
  - Significant changes to SMS ACS routines
  - Significant changes to SMS constructs, especially target thresholds of storage groups.

Note to New User: All issues apply, with the obvious exception of the long-term running of MMC.

If You Do Not Perform Quick-Start

If you are either an existing user of MMC or a new user of MMC and you do *not* want to perform Quick-Start, then perform the following steps:

Step	Action
1	Use the ASXCALC CLIST to compute database size requirements.
2	Use ALLOCMMC in ASX.CNTL to allocate the new MMC database. (Make sure you know the space requirements for the MMC database before doing this step.)
3	Format the new database by running the job in member AMMFMTDB, also in ASX.CNTL.

Step 4: Tailor the ISPF Options

Specifying Values for the Skeleton JCL

You should edit and tailor the ASXPARMS member of the ASX.ASX28.SLIB data set. Comments in this member instruct you to specify various values for the skeleton JCL used by the CA-ASTEX batch components. In particular, you need to specify correct data set names for certain load libraries, UNIT names for work data sets, space allocation information for SORTWK*nn* data sets, and other JCL parameter values.

Note: Users can make temporary changes to the skeleton JCL before submitting CA-ASTEX batch jobs. They need only to specify YES at the following prompt on a batch submittal screen:

```
Edit the JCL before submitting the job:  YES
```

Each user can make permanent changes to the JCL by modifying default values via the CA-ASTEX "User Defaults" (option A1.1 on the Primary Menu).

Establishing ISPF Access to CA-ASTEX

To establish access to CA-ASTEX from ISPF, you must allocate the CA-ASTEX ISPF data sets, and provide access to the ISPF features.

You allocate the CA-ASTEX ISPF data sets by creating a CLIST with LIBDEF statements that point to the CA-ASTEX ISPF data sets. A sample CLIST, shown here, is in the ASXCLIST member of ASX.CNTL:

```
PROC 0
   CONTROL NOMSG NOFLUSH
   ATTN GOTO CLEANUP
   ALLOC F(ASXDUMMY) DUMMY
   SET CC = &LASTCC
   CONTROL MSG

   IF &CC > 0 THEN DO
      ISPEXEC LIBDEF ISPMLIB DATASET ID('ASX.ASX28.MLIB')
      ISPEXEC SETMSG MSG(ASX547)
      EXIT
   END

   ELSE DO
     CONTROL NOMSG
     FREE F(CTRANS)
     CONTROL MSG
     ALLOC F(CTRANS) DA('ASX.ASX28.LOAD') SHR
     ISPEXEC LIBDEF ISPPLIB DATASET ID('ASX.ASX28.PLIB')
     ISPEXEC LIBDEF ISPSLIB DATASET ID('ASX.ASX28.SLIB')
     ISPEXEC LIBDEF ISPMLIB DATASET ID('ASX.ASX28.MLIB')
     ISPEXEC LIBDEF ISPLLIB DATASET ID('ASX.ASX28.LOAD')
     ISPEXEC SELECT PGM(ASXTSOM) NEWAPPL(ASX$)  PASSLIB
   END

CLEANUP: +
   ATTN OFF
   FREE F(ASXDUMMY)
   FREE F(CTRANS)
   ISPEXEC LIBDEF ISPPLIB
   ISPEXEC LIBDEF ISPSLIB
   ISPEXEC LIBDEF ISPMLIB
   ISPEXEC LIBDEF ISPLLIB
   EXIT
```

Note: Computer Associates recommends that you use the CLIST member above. This CLIST prevents two CA-ASTEX monitor sessions from being started in the same TSO session. Without this CLIST, if someone attempts to start a second monitor session, the TSO user ID hangs.

Once the ISPF data sets are allocated, you can provide ISPF access to the CA-ASTEX Primary Menu (and in turn to all CA-ASTEX screens) through a command, the ISPF menu, or both.

- To provide access to CA-ASTEX via a %ASTEX command issued from ISPF option 6 or a screen command line, place your CLIST in a member named "ASTEX."

- To provide access to CA-ASTEX from the ISPF menu, you must add two lines of code on your ISPF/PDF Primary Option Menu screen. If you allocated your data sets by creating a CLIST, change the line of code indicated by the arrow below to the following:

  ```
  A,'CMD(%ASTEX)'
  ```

Note: Using similar code, you could create a CA-ASTEX selection on a secondary option menu instead:

```
%----------------------- ISPF/PDF PRIMARY OPTION MENU ---------------------
%OPTION ===> ZCMD
%                                                              +USERID   -
% 0 +ISPF PARMS  - Specify terminal and user parameters        +TIME     -
% 1 +BROWSE      - Display source data or output listings      +TERMINAL -
% 2 +EDIT        - Create or change source data                +PF KEYS  -
% 3 +UTILITIES   - Perform utility functions
% 4 +FOREGROUND  - Invoke language processors in foreground
% 5 +BATCH       - Submit job for language processing
% 6 +COMMAND     - Enter TSO command or CLIST
% 7 +DIALOG TEST - Perform dialog testing
% A +CA-ASTEX    - Automated Storage Performance Expert
% C +CHANGES     - Display summary of changes for this release
% T +TUTORIAL    - Display information about ISPF/PDF
% X +EXIT        - Terminate ISPF using log and list defaults
%
+Enter%END+command to terminate ISPF.
%
)INIT
  .HELP = ISR00003
  &ZPRIM = YES        /* ALWAYS A PRIMARY OPTION MENU     */
  &ZHTOP = ISR00003   /* TUTORIAL TABLE OF CONTENTS       */
  &ZHINDEX = ISR91000 /* TUTORIAL INDEX - 1ST PAGE        */
)PROC
  &ZSEL = TRANS( TRUNC (&ZCMD,.)
                0,'PANEL(ISPOPTA)'
                1,'PGM(ISRBRO)'
                2,'PGM(ISREDIT)'
                3,'PANEL(ISRUTIL)'
                4,'PANEL(ISRFPA)'
                5,'PGM(ISRJB1) PARM(ISRJPA) NOCHECK'
                6,'PGM(ISRPTC)'
                7,'PGM(ISRYXDR) NOCHECK'
     ------->   A,'CMD(%ASTEX)'
                C,'PGM(ISPTUTOR) PARM(ISR00005)'
                T,'PGM(ISPTUTOR) PARM(ISR00000)'
              ' ',' '
                X,'EXIT'
                *,'?' )
  &ZTRAIL = .TRAIL
)END
```

This step is complete after you have allocated the CA-ASTEX ISPF data sets and created a CA-ASTEX ISPF menu option.

Step 5: Tailor and Submit the ASXFMTDB Job to Initialize the PMC Databases

Before you start the CA-ASTEX installation procedure, you must initialize the appropriate databases. The ASXFMTDB member of the ASX.CNTL data set contains instructions on how to do this.

Step 6: (Optional) Install the ARCRDEXT Exit

If you are using DFHSM, you can use the CA-ASTEX Allocation Director to optimize any non-directed DFHSM recalls of data sets that are not managed by SMS. To do this, you must install the CA-ASTEX ARCRDEXT exit on each system that runs CA-ASTEX and performs DFHSM recalls for non-SMS managed data sets.

Follow these steps to install the CA-ASTEX ARCRDEXT exit:

Step	Action

1	Determine whether ARCRDEXT is already installed at your site. Look for the following statements in the DFHSM parameter library:
	`| SETSYS -`
	`| EXITON(ARCRDEXT)`
	If these statements do exist, a user-written ARCRDEXT exit is currently installed at your site. Proceed to Step 2.
	If these statements do not exist, you do not currently have an ARCRDEXT exit installed at your site. Proceed to Step 3.

2	Rename your current ARCRDEXT exit as a new exit called ARCRDUSR. The CA-ASTEX ARCRDEXT exit then passes control to your user-written exit (now named ARCRDUSR) when a non-SMS HSM recall allocation takes place.
3	Link-edit the CA-ASTEX ARCRDEXT exit into a linklist data set. This load module can be found in your CA-ASTEX authorized library. Sample JCL for performing the link-edit can be found in the LNKARCRD member of the ASX.CNTL data set.

4	Issue the following DFHSM commands to ensure that DFHSM uses the newly-linked CA-ASTEX exit:
	`SETSYS EXITOFF(ARCRDEXT)`
	`SETSYS EXITON(ARCRDEXT)`

Step 7: (Optional) Set Up Connectivity With BrightStor DASD Manager Option for Unicenter and CA-ASTEX

Perform this step if you want to use the BrightStor DASD Manager Option for Unicenter and CA-ASTEX to view CA-ASTEX alerts and historical reports. The following diagram illustrates how the components installed in this step are connected:

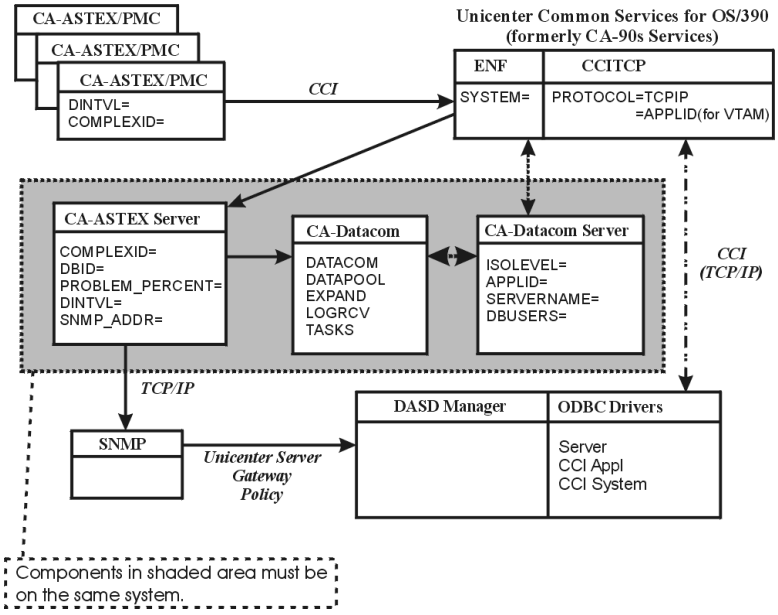

The following steps must be completed before CA-ASTEX is able to produce the storage hierarchy map and send alerts to its Unicenter server and the BrightStor DASD Manager Option for Unicenter and CA-ASTEX.

Step	Action
1.	Install the IBM TCP/IP with SNMP component. See the TCP/IP documentation for instructions on installing TCP/IP.
2.	Install the following components from the Unicenter Common Services for OS/390 tape shipped with CA-ASTEX 2.8: ■ CAIENF ■ CAICCI Refer to the Unicenter Common Services for OS/390documentation for instructions. Use the following settings: a. For CCI, `PROTOCOL=TCPIP` `PROTOCOL=APPL` (*for VTAM*) b. For ENF,

Step	Action
	SYSTEM=*sysname*
3.	Install either CA-Datacom/DB 4.0 or CA-Datacom/AD, version 9.0, or higher. CA-Datacom/AD is supplied on a distribution tape shipped with CA-ASTEX 2.8. You must install CAIENF and CAICCI (listed in Step 2) prior to installing CA-Datacom/AD. Refer to the CA-Datacom documentation for instructions on installing CA-Datacom/AD. If CA-Datacom is already installed, the databases needed by CA-ASTEX can be installed on the existing CA-Datacom installation.
4.	Install the CA-Datacom Server, version 4.0 or higher. The CA-Datacom Server is supplied on a distribution tape shipped with CA-ASTEX. Refer to the CA-Datacom documentation for instructions on installing the CA-Datacom Server. The following parameters must be set for the CA-Datacom Server to accommodate CA-ASTEX: ISOLEVEL=U APPLID=*applid* SERVERNAME=*servername* DBUSERS=*X* (where *X* is the number of users to be viewing reports concurrently.) AUTHID=SYSUSR

Step	Action
5.	Allocate the CA-ASTEX Server Datacom Alert Database and define tables to CA-Datacom.

a. Use the following formulas to calculate the number of tracks needed for the data and index components of the database:

L = number of LCUs in the complex
V = number of volumes in the complex
I = interval duration in minutes
D = number of days to keep in database
F = device factor. 3390's use 622, 3380's use 519

Tracks for data component = `((L+V)/F) * ((1440/I) * D)`
Tracks for index component = *tracks for data component* * `.35`

For example, given the following values

L = 28

V = 500

I = 10

D = 3

F = 622 (3390 device)

The calculation for the data component is:

((28 + 500) / 622) * ((1440/10) * 3)

(528/622) * (144 * 3)

.85 * 432

367.2 tracks or 368 tracks rounded

And the calculation for the index component is:

368 * .35

128.8 tracks or 129 tracks rounded

b. Modify and submit the following JCL:

ASX.CNTL(ALLOCDB), if using CA-Datacom/DB

ASX.CNTL(ALLOCAD), if using CA-Datacom/AD

Note: Before submitting this job, make sure the CA-Datacom multi-user facility (DBMUF) is running.

Instructions for modifying the JCL are provided in the member. These jobs allocate the data sets for the database and run CA-Datacom utilities that prepare the data sets for use.

Step	Action
6.	Add the IP addresses of your Unicenter server to the hlq.SNMPTRAP.DEST data set for each Unicenter server that is to receive alerts from the CA-ASTEX Server. This parameter is part of the TCP/IP SNMP component, and can be defined in either a sequential data set, or in a member of a partitioned data set. The format is: *nnn.nnn.nnn.nnn* UDP
7.	Use the supplied procedure contained in member PROCSRV of the ASX.CNTL data set, to start the CA-ASTEX server. The CA-ASTEX server must run on the same system running CA-DATACOM/AD or CA-DATACOM/DB. It does not need to run on the same system as the IBM SNMP agent.
8.	Examine the supplied parms member ASXPARMS. Use the diagram at the beginning of this step and the following tables to make sure the proper parameters correspond on the various components. You **must** specify a value for parameter SNMP_ADDR=.

CA-ASTEX Component Settings

Component	Setup	Comment
Unicenter Common Services for OS/390 ENF	SYSTEM=*sysname*	
CCI/TCP	PROTOCOL=APPLID(*for VTAM)*	The VTAM APPLID, for cross-OS/390 systems only.
	PROTOCOL=TCPIP	for Workstation communication
CA-Datacom	DATACOM=DTCMSRVR	Identifies Datacom server to MUF
	DATAPOOL *dataln, datano*	*dataln*>=4K
	EXPAND *length*, [*number*]	*length*>=8K
	LOGRCV *entry*	*entry* =NO or NEVER
	TASKS *number*, *size*	*size*>= 16K *number*= 1 + value of DBUSERS parameter on the CA-Datacom server * *If other products are using the MUF in addition to CA-ASTEX, you need to add the recommended number for those products also.

Component	Setup	Comment
CA-Datacom Server	ISOLEVEL=U	No locks required, no changes are allowed
	APPLID=*applid*	
	SERVERNAME=*servername*	
	DBUSERS=*X*	*X*= the number of users to be viewing reports concurrently.
	AUTHID=SYSUSR	
CA-ASTEX (PMC)	ALERTS=YES	Needs to be YES to process alert data.
	COMPLEXID=*compname*	Needs to match COMPLEXID for the CA-ASTEX Server
CA-ASTEX Server	COMPLEXID=*compname*	Needs to match COMPLEXID for CA-ASTEX PMC.
	PROBLEM_PERCENT=	
	SNMP_ADDR=	The IP name or address of the machine running the SNMP agent.
	DINTVL=	Must match DITNVL on **all** PMCs in the complex.
with Datacom	DBID=	/AD uses 620 or default /DB uses value set in ALLOCDB install job

See the chapter "Tailoring CA-ASTEX Parameters" for details on defining CA-ASTEX parameters.

Parameter Correspondences Across Components

This Parameter...	on this component...	corresponds to this parameter...	on this component.
APPLID	CA-Datacom Server	CCI Appl	CA-Datacom ODBC Setup (Workstation)
COMPLEXID	CA-ASTEX PMC	COMPLEXID	CA-ASTEX Server
DBID	CA-Datacom/AD CA-Datacom/DB	DBID	CA-Datacom Server
DINTVL	CA-ASTEX PMC	DINTVL	CA-ASTEX Server
SERVERNAME	CA-Datacom Server	SERVER	CA-Datacom ODBC Setup (Workstation)
SYSTEM	CAIENF	CCI System	CA-Datacom ODBC Setup (Workstation)

Step 8: (Optional) Set Up Security for the Command Interface

You do not have to perform this step if you set up security for the Command Interface in an earlier version of CA-ASTEX.

You can prevent unauthorized users from issuing CA-ASTEX operator commands by following these steps:

Step	Action
1	If your security package is not RACF, make sure it is set up to use SAF. CA-ASTEX uses the RACROUTE macro and the system authorization facility (SAF).
2	Your security administrator must add a generic command class name that is also defined as an active class. The name of this command class should be @ASXCMD. Your security administrator also must add a resource name, MODIFY, to this class. Authorized users should be given READ access; all others should be given an access of NONE.
3	Specify SECURE=YES in the DEFINE PARMS section. If your security package is not set up to use SAF, you must specify SECURE=NO.
4	Add the new @ASXCMD resource class to the class descriptor table and the router table. Make sure that the resource name (MODIFY) has a maximum length of eight characters. In order for the changes to take place, you must re-IPL.

See the CA-ASTEX *User Guide* chapter "Understanding the CA-ASTEX Facilities" for a description of the Command Interface and how to use it.

Step 9: Tailor the PMC Start-up Procedure

Member PROCPMC in the ASX.CNTL data set contains a sample start-up procedure for PMC. Tailor the procedure to suit your needs.

1. Follow the customization instructions contained in PROCPMC, paying particular attention to the following statements:

   - **The PROC statement at the beginning of the procedure**

     This statement contains default values for the REUSE, LOG, TDB, and SCOPE parameters, which are invoked when CA-ASTEX is started:

     ```
     //PROCPMC   PROC  LOG=YES,OPTS=(D0),REUSE=YES,SCOPE=PMC,TDB=tdb name
     ```

Parameter	Explanation
LOG=YES	This parameter indicates that you want the CA-ASTEX PMC message log allocated to a SYSOUT data set, class A.
OPTS=(D0)	This parameter indicates which CA-ASTEX options your site is authorized to use. Options must be enclosed in parentheses. Separate multiple options with commas.
REUSE=YES	This parameter indicates that, if possible, you want PMC to reuse common storage left from a previous execution. You should *always* use the default value of YES. If PMC cannot reuse storage, the value is changed from YES to NO.
SCOPE=PMC	This parameter indicates that this start-up procedure is for PMC.
TDB=*tdb name*	Specify the name of the PMC Trend Database that you allocated.

These statements contain the names of actual data sets you are using:

DD statement	Data set it identifies
//ASXPARM DD	CA-ASTEX parameter data set (Default: ASXPARM member of ASX.CNTL)
//SYSTERM DD	Output data set for C++ PRINTF and SYSOUT
//ASXIDB DD	Optional Interval Database containing summary and detail data or a file containing only summary data. (Default: None)
//ASXIDBD DD	Optional Detail Interval Database file containing only detail data. (Default: None)
//SNAP0001 DD	Output data set where snap dump is directed (Default: None)
//ASXTDSN DD	Optional user-defined data set for trace records (Default: None)

2. Create a new member named "ASXPMC" in your system procedure library, and copy the contents of member PROCPMC in the ASX.CNTL data set to this PROCLIB data set.

Step 10: Tailor the MMC Start-up Procedure

This step is necessary only if you are running MMC.

Member PROCMMC in the ASX.CNTL data set contains a sample CA-ASTEX start-up procedure for MMC, which you should tailor to suit your needs.

1. Follow the customization instructions contained in PROCMMC, paying particular attention to the following statements:

   - **The PROC statement at the beginning of the procedure**

     This statement contains default values for the LOG and SCOPE parameters, which are invoked when CA-ASTEX is started:

     ```
     //PROCMMC    PROC    LOG=YES,OPTS=(FT),PROG=ASXMSTR,SCOPE=MMC
     ```

Parameter	Explanation
LOG=YES	This parameter indicates that you want a CA-ASTEX message log allocated to a SYSOUT data set, class A.
OPTS=(FT)	This parameter indicates which CA-ASTEX options your site is authorized to use. Options must be enclosed in parentheses.
PROG=ASXMSTR	This parameter invokes various batch utilities and report programs. Do not modify this.
SCOPE=MMC	This parameter indicates that this start-up procedure is for MMC.

   - **The ASXPARM and HSMPARM DD statements**

     These statements contain the names of actual data sets you are using:

DD statement	Data set it identifies
//ASXPARM DD	CA-ASTEX parameter data set (Default: ASXPARM member of ASX.CNTL)
//HSMPARM DD	DFHSM ARCCMD*xx* member for the system on which CA-ASTEX is to be started. This statement is not needed if you have coded the HSMPARM= option in the ASXPARM member.

2. Specify performance characteristics for the PROCMMC started task that is consistent with the performance characteristics of a high-priority batch job at your installation.

3. Create a new member named "ASXMMC" in your system procedure library, and copy the contents of member PROCMMC in the ASX.CNTL data set to this PROCLIB data set.

Using the Quick-Start Facility (MMC Only)

The Quick-Start facility allows you to "prime" the Migration Manager database with historical HSM activity data and observed access patterns for active data sets. The historical data allows you to establish a baseline of previous HSM activity to better compute Migration Manager benefits, such as a reduction in data set recalls. The access pattern data allows the Migration Manager to make intelligent movement recommendations sooner, rather than having to wait several cycles to "learn" about the data.

This facility is optional, but *we strongly recommend* that you take the time to load historical HSM activity data before commencing MMC operations.

You benefit from Quick-Start processing if:

- You are just implementing the Migration Manager and have never run a daily update cycle. You can use the Quick-Start facility to prime your database with historical data before you run your first daily cycle. Note the instructions for *new users* in the section Checklist for All Users in this chapter.
- You implemented Migration Manager a short time ago but have not run many daily cycles. You can use the Quick-Start facility to add historical data to your database for the time period immediately before you started running the Migration Manager.

Reasons for Using Quick-Start

Here are some reasons why you should use Quick-Start:

- Batch Reporting of HSM activity is greatly enhanced, allowing you to better see long-term trends on a variety of data measurements.
- The effectiveness of the Migration Manager is evident sooner, arising from the fact that MMC detects certain access patterns directly from the HSM activity of the past.
- Quick-Start is a "one-time" activity that mostly involves a lot of tape mount activity. The amount of CPU-time is minimal. Your own time is not affected either, because you can do other activities while the Quick-Start jobs are running. Depending on how you set up the JCL and whether you have robotic tape drives, the overall Quick-Start period may take a few days, but this time should not affect your own workday except very indirectly.

Processing Overview

The Migration Manager started task traps HSM Activity records (FSR records) in an IEFU83 exit and then incorporates those records into the database. This gives you a good summary of the work being done each day by HSM requests, Migration Manager requests, and user-generated requests. In a similar fashion, the Quick-Start facility will:

- Read your historical SMF data sets to extract FSR records and RMF type 74 records.
- Use your RMF records to obtain storage group and capacity values for each volume.
- Tabulate HSM activity and error counts by volume for each day.
- Identify access patterns for active data sets represented by the FSR data.
- Merge the FSR and RMF data.
- Create intervals in the database to represent HSM activity for each day.
- Create a current interval containing access patterns for active data sets.

Checklist for All Users

Follow this checklist of steps to accomplish Quick-Start. Each utility related to Quick-Start is executed by a sample job stream contained in ASX.CNTL.

Step 1: Plan for Quick-Start

Follow these steps to plan for Quick-Start:

1. The choice you make for the parameter PREFIX has a major impact on the processing of both the Quick-Start series of processes and the MMC address space itself. See the chapter "MMC: Planning Before Customizing," especially the section "Importance of the PREFIX Parameter."
2. Find out where your historical SMF data is stored. In most cases, there is a naming convention for dumped SMF data, and it is typically organized in a generation data group (GDG). For discussion purposes and a useful example, consider an HSM complex with four OS/390 images and assume that HSM migrates and recalls can occur on each system.

The following table contains sample SMF data structures for Quick-Start.

system 1	PROD.SMFDUMP.S001.MONTHLY PROD.SMFDUMP.S001.WEEKLY PROD.SMFDUMP.S001.DAILY
system 2	PROD.SMFDUMP.S002.MONTHLY PROD.SMFDUMP.S002.WEEKLY PROD.SMFDUMP.S002.DAILY
system 3	PROD.SMFDUMP.S003.MONTHLY PROD.SMFDUMP.S003.WEEKLY PROD.SMFDUMP.S003.DAILY
system 4	PROD.SMFDUMP.S004.MONTHLY PROD.SMFDUMP.S004.WEEKLY PROD.SMFDUMP.S004.DAILY

Step 2: Perform Phase 1 Processing (All Users)

Run member AMMQUIC1 to perform Phase 1 processing. This job stream is designed to:

- Read your SMF data.
- Reduce the FSR and RMF records to keep only the necessary items.
- Save the data in an accumulation data set.

You can run this job stream multiple times, using different sources of input SMF data for each run. The order the data is fed into the utility is not important, as later processing orders the data into the correct time sequence. Because the basis of this data is to establish a historical baseline of HSM activity, we suggest that you use at least three to six months of historical SMF data. Less than that is not enough for comparison purposes, while more than that may reflect configuration changes in your environment that skew the results.

For example, in the previous table containing SMF data structures for Quick-Start, various cataloged data sets contained dumped SMF data. Assuming you are starting Quick-Start processing in the middle of the month, you should input the following:

system 1	PROD.SMFDUMP.S001.MONTHLY *all generations* PROD.SMFDUMP.S001.WEEKLY(0) PROD.SMFDUMP.S001.WEEKLY(-1) PROD.SMFDUMP.S001.DAILY(0) PROD.SMFDUMP.S001.DAILY(-1) PROD.SMFDUMP.S001.DAILY(-2)
system 2	Repeat same input sequence for SMF data from this SYSID.
system 3	Repeat same input sequence for SMF data from this SYSID.
system 4	Repeat same input sequence for SMF data from this SYSID.

In this example scenario, the latest "monthly" generation will not contain the period-to-date data for the month you are currently in; therefore, the Quick-Start data is more complete if you include one or two weeks of SMF activity and *n* days of activity, too.

There is a lot of tape mounting activity for these jobs. Make sure you anticipate this with proper JOB parameters for your installation, paying particular attention to resource controlling functions such as TIME and CLASS.

If you get a correctable error, you should feel free to rerun the job. There is likely to be duplicate data, but this is filtered out later, so there is no need to worry about that or to start over with an empty accumulation data set.

It is unlikely that an out-of-space error (Sx37 abend) will occur during AMMQUIC1 processing. But if an error does occur, perform the following steps:

Step	Action
1	Allocate a larger data set with the same attributes as the data set in error.
2	Copy the old data set into the larger data set.
3	Delete the old data set.
4	Rename the larger data set to the same name as the old data set.

Use the report(s) produced by this utility to make sure your history is complete (without gaps) before proceeding to the next step. If you have "minor gaps," you can still proceed; the point is to be as reasonably all-inclusive as possible.

Note for Existing MMC 2.7 Users: Assume you have been running MMC 2.5 for about three months. If your "oldest date" in the MMC database is, for example, 20Mar00, then you should strongly consider adding some Quick-Start data to your database to provide a long-term view of summary data related to HSM activity in your installation. Assume you decide to collect data for the period including calendar year-end 2000, so you begin this phase of Quick-Start processing with data having date stamps starting on, for example, 15Dec00.

Do not be concerned about the overlap of Quick-Start processing with data already in the database. That is, you can safely include data for the month of March (2000) in this example. Later Quick-Start processing (Phase 3 discussed in the section "Checklist of All Users/New Users Follow These Steps") recognizes that Quick-Start data overlaps with data already in the MMC database. If this occurs, Quick-Start ignores the Quick-Start data and keeps the MMC database data because it is more comprehensive. So, while there is no concern about duplicate data, you should try to prevent extreme cases of overlap like this because it wastes resources. Quick-Start data is also ignored if it has a date later than the most current interval in the database.

Step 3: Perform Phase 2 Processing (All Users)

Run member AMMQUIC2 to perform Phase 2 processing. This job stream is designed to:

- Read your accumulation data set.
- Tabulate the historical data by day and the access data by data set.
- Create the necessary objects that are written to the database.
- Write those objects to a permanent file used by the next job stream.
- Delete the accumulation data set.

A summary report is produced that estimates the size of the DBMAIN and DBFILE components of the database, computed based on the amount of data processed. This size is the minimum you should use to allocate your database. Using an amount smaller than this may cause some of your historical statistics to "roll off" the database before they can be used in your comparisons.

Warning! *Because the job deletes the accumulation data set, you cannot go back and run more data through AMMQUIC1 after you have run AMMQUIC2. Therefore, because extensive tape mounts were required, we recommend that you take an independent backup of the* prefix.* *data sets now (in case you want to extend AMMQUIC1 later).*

Step 4: Compute Database Size Requirements

Use the ASXCALC CLIST to compute the space requirements for the database. You should also use the summary output report from AMMQUIC2.

Step 5: Allocate the New MMC Database

Use ALLOCMMC in ASX.CNTL to allocate the new MMC database. (Make sure you know the space requirements for the MMC database before doing this step.)

Step 6: Format the New MMC Database

Format the new database by running the job in member AMMFMTDB, also in ASX.CNTL.

At this point, there are separate paths in the checklist, depending on whether you are a new user or a converting user.

New Users Follow These Steps

If you are a new user of Migration Manager, you have no database to convert; you have only Quick-Start data. Perform the following steps:

Step	Action
1	Run member AMMQUIC3 to perform Phase 3 processing. This job stream is designed to read the permanent file created by the previous step, and create intervals in your new database to represent your data. An interval is created for each date for which HSM activity statistics are found. Access patterns for the active data sets found in your SMF data are stored in the most recent interval, much like normal daily cycle processing. You can immediately run some of the Trend reports on the database to view the trend of HSM activity over the "Quick-Start" period of time. When you run a daily space management cycle for the first time the Migration Manager has immediate knowledge about the access patterns of some of your more active data sets from the Quick-Start data. Thus, besides enabling quicker identification of data set access patterns for the quicker implementation of MOVEMODE=ON, Quick-Start processing gives you an unbroken chain of summarized HSM activity history.

2 Perform the remaining steps:

- At this point, you have a database primed with summarized history of HSM activity and some data set level access patterns. You should proceed toward getting the MMC address space operational now.
- While you are establishing PROCMMC as a started task address space to be run continuously on all systems, you can also get some batch reporting accomplished in parallel to these setup activities. You can also start using the ISPF dialogs, but be aware that there is only very limited data available for inspection with only Quick-Start intervals in the database.

Starting and Checking CA-ASTEX PMC

The following instructions tell you how to start the CA-ASTEX Performance Management Components and check to see whether they are working properly:

1. Start the CA-ASTEX Performance Management Components by entering the following command:

   ```
   S ASXPMC
   ```

   This command activates the PROCPMC start-up procedure and initializes the CA-ASTEX Performance Management Components' data sets. When initialization is complete, you receive a message that looks like this:

   ```
   ASX001I - CA-ASTEX v.r SERVICEPACK NNNN PERFORMANCE MANAGEMENT COMPONENTS ARE
   ACTIVE
   ```

2. Check the "status" of the CA-ASTEX Performance Management Components by entering the following command:

   ```
   F ASXPMC,STAT
   ```

   This command asks the system to confirm whether CA-ASTEX is "active" or "inactive." If the CA-ASTEX Performance Management Components are active, the system replies with a message that looks like this:

   ```
   ASX001I - CA-ASTEX v.r SERVICEPACK NNNN PERFORMANCE MANAGEMENT COMPONENTS ARE
   ACTIVE
   ```

3. Go to the CA-ASTEX Primary Menu by doing one of the following:
   - Select CA-ASTEX from an ISPF option menu.
   - Issue a %CA-ASTEX command.
4. In the Primary Menu, select the option you want to execute. For more information about these options and how to use them, please refer to the following guides and their topics:

- CA-ASTEX *Storage Manager User Guide*, chapter "Understanding the Storage Manager Component of CA-ASTEX," section "The Storage Manager Facilities."
- CA-ASTEX *Cache Manager User Guide*, chapter "Understanding the Cache Manager Component of CA-ASTEX," section "How You Can Use the Cache Manager."
- CA-ASTEX *DASD Manager User Guide*, chapter "Understanding the DASD Manager Component of CA-ASTEX," section "How You Can Use the DASD Manager."

Starting and Checking the CA-ASTEX Server

Follow the steps below to start the CA-ASTEX server and to check to see that it is working properly.

1. Create a new member named "ASXSERV" in your system procedure library, and copy the contents of member PROCSRV in the ASX.CNTL data set to this PROCLIB data set.
2. Start the CA-ASTEX Server by entering the following command at the MVS console:

   ```
   S ASXSERV
   ```

   The command activates the ASXSERV startup procedure. When initialization is complete, the following message is displayed:

   ```
   ASX301I - CA-ASTEX v.r Servicepack nnnn is active for complex cccccccc
   ```

   A sample startup JCL is located in ASX.CNTL member PROCSRV.
3. Stop the CA-ASTEX Server by typing a purge command:

   ```
   P ASXSERV
   ```
4. To display the status of the CA-ASTEX Server, enter the following command at the OS/390 console:

   ```
   F ASXSERV,STAT
   ```

   This command returns the following messages:

   ```
   ASX301I - CA-ASTEX v.r Servicepack nnnn is active for complex cccccccc
   ASX318I - Task taskname status
   ASX318I - Task taskname status
   ```

Starting and Checking CA-ASTEX MMC

The following instructions tell you how to start the CA-ASTEX Migration Manager Component and check to see if it is working properly:

1. Start CA-ASTEX MMC by entering the following command:

   ```
   S ASXMMC
   ```

   This command activates the ASXPROCM start-up procedure and initializes the CA-ASTEX MMC data sets. When initialization is complete, you receive a message that looks like this:

   ```
   AMM1000I - CA-ASTEX v.r Servicepack nnnn  Migration Manager Component is Active
   ```

2. Check the status of CA-ASTEX MMC by entering the following command:

   ```
   F ASXMMC,STAT
   ```

   This command asks the system to confirm whether CA-ASTEX MMC is active or inactive. If it is active, the system replies with a message that looks like this:

   ```
   AMM1000I - CA-ASTEX v.r Servicepack nnnn Migration Manager Component is Active
   ```

3. Go to the CA-ASTEX Primary Menu by doing one of the following:

   - Select CA-ASTEX from an ISPF option menu.
   - Issue a %ASTEX command.

   Note: Unless you used Quick-Start data to build your database, there is no data to display until the Migration Manager has completed at least one daily inventory cycle.

4. In the Primary Menu, select the option you want to execute. For more information about these options and how to use them, please refer to the following guide and topics:

   CA-ASTEX *Migration Manager User Guide*, in the section "How to Access the Migration Manager" in the chapter "Introducing the Storage Manager Component of CA-ASTEX."

Complete the Processing of the MMC Daily Cycle Using RESTART=FORCE

The *daily cycle* is the address space processing that builds and maintains the database and chooses data sets for migration. The *first daily cycle after install is critical!* Completion of the first daily cycle (and the second daily cycle) represents a *functional verification procedure*.

The daily cycle is triggered by the START=*hhmm* parameter. It will not start until that time of day.

To verify installation and implementation, you should not wait until the START time specified in ASXPARM (that time should be chosen relative to when the HSM daily cycle starts).

Rather than wait for the START time, we recommend that you issue the modify command to trigger the daily cycle as soon as you are ready, after starting the address space. Example syntax:

```
F ASXMMC,RESTART=FORCE
```

This way, you can observe the daily cycle processing and examine the JOURNAL file as it is running, and especially after it is finished. You should see informational messages about the processing phases completing normally. If you see warning or error messages, you can go immediately to the CA-ASTEX *Message Guide* for more information. You can call Computer Associates Technical Support if resolution of the problem is not clear.

Do not bother using the ISPF dialogs for CA-ASTEX MMC until the first daily cycle has completed successfully. Until it does, there is no data to display. Browse the JOURNAL DD for the ASXMMC address space for the message:

```
AMM1041I Daily Space Management Ended with No Errors.
```

Some products used for browsing SPOOL for active jobs buffer the most recent output messages, so be aware of this possibility.

Assuming the first daily cycle running as a result of RESTART=FORCE is successful, *what happens when the actual START time is reached*? Will the daily cycle start again? The answer is no, if the most recently completed daily cycle ran less than 12 hours previously; if the previous daily started more than 12 hours before the triggered START time, then the daily cycle starts. Look for message AMM1050I in the log; it indicates when the next daily cycle runs.

Here is an example:

- On 11 April 01 at 13:05, you issue F ASXMMC,RESTART=FORCE.
- On 11 April 01 at 22:00, the START=2200 triggers a timer interrupt, which tests if the daily should start.
- Because the previous daily started fewer than 12 hours previously, the daily is skipped.
- On 12 April 01 at 22:00, the timer interrupt is triggered and the daily cycle starts.
- Subsequent days are processed in the same manner, and the length of the intervals should become 24 hours.

Performance Considerations on Test Systems

Implementation of program products such as MMC typically starts with installation and operation on a "test" system. Whether this is a guest OS/390 running under VM or whether it is an LPAR, we have found that test systems are rarely given substantial access to CPU (processor) resources. For example, a test system may be granted only 5% or 10% of the total processing power of a processor complex.

MMC, however, has "production work" to do even on a test system, and if not given proper access to the CPU, the run time of the address space is longer (on the test system) than when it is placed into production.

This production work involves scanning all VTOCs and VVDSes, rearranging and sorting this raw data, scanning the complete MCDS, and building similar intermediate files. Subsequent phases involve building DSNAME node tables in extended memory (above 16m line) and more sorting of the data back into the physical sequence(s) for the final database build process.

If, for example, this production work is performed from a "test" LPAR on a large shared DASD configuration (many online VOLUMES in the domain of the Migration Manager), it would be quite natural and expected to see relatively long run times via the time stamps of messages in the JOURNAL data set.

Assuming the DASD complex to be examined and managed remains the same, this elapsed time should shrink dramatically when moved to a production LPAR, as discussed in the next section.

Performance Considerations on Production Systems

Migration Manager issues a SYSEVENT to make the address space non-swappable for the duration of the daily cycle. After the daily cycle, the address space is again marked swappable as it reverts to collecting FSR (HSM activity) data only.

Because the entire daily cycle can generate very large quantities of I/O requests and CPU utilization, it is very important that Migration Manager be run in a proper performance group.

Performance Group Should Be For "Hot Batch"

Favored batch, or "hot" batch, is the best descriptor of the performance required by Migration Manager. During the daily cycle, MMC is behaving like a batch job with large data requirements. Because the ultimate target is generation of an HSM workload to groom volumes, the profile of this address space is also related to system availability (adequate free space), and hence it should be favored.

Note: It is vitally important that the priority be lower than online systems such as CICS.

During the daily cycle, monitors show that MMC generates high CPU utilization and lots of EXCPs. This is normal and expected, but it is important that a workload like this does not interfere with the response time requirements of online production systems. In simple terms, MMC needs to be set up so that it gets all the "horsepower" that is available but relinquishes control to jobs and tasks that need higher priority whenever they need to be dispatched.

Updating SYS1.PARMLIB (IEAICSxx) and Possibly IEAIPSxx

To affect these settings, IEAICS*xx* should contain an entry for the ASXMMC started task job name that assigns it a "hot batch" PERFORM group.

If a facility for "hot batch" is not available, or if you want to make a separate performance group for MMC only, then you need to update IEAIPS*xx* to define the performance group itself. The IEAICS*xx* entry must refer to a defined performance group.

MVS Version 5.x Users of the Workload Manager

The same principles outlined above apply in MVS/ESA Version 5.1 and above. When the Workload Manager is configured to run in goal mode, you must ensure that the equivalent setup is made for the service class assigned to Migration Manager.

Choice of START Time Can Influence Elapsed Time

Whether running on a test LPAR or running on a production system, the choice you make for the Migration Manager START time can have some effect on the duration of the daily cycle. Obviously, if you can select a time when there is relatively little contention for the resources needed by MMC (that is, CPU, access to MCDS, access to *prefix.ddname*, and so on), there are better durations than where there is contention for these resources.

In general, however, it is more important to choose a START time that is 2 to 4 hours before the time you have already chosen for the HSM daily space management (assuming that the HSM choice was as well-made as possible, given other operational considerations).

Excessive MMC Migration Recommendations

Remember that MMC obeys the TARGET threshold inherited from SMS or HSM and tries to manage to that target without regard for eligibility. If you give MMC a very low target based on HSM philosophies (for example, a target less than 30% or something unrealistic like 1%), then MMC is forced to find sufficient candidate data sets to achieve that target.

Excessive migration requests can lead to many problems and should be avoided, especially when unintended by failing to use the TARGET parameter of DEFINE SGROUP to override the very low targets of HSM.

Even with MOVEMODE=OFF, where the migration requests exist only as recommendations, excessive amounts can cause elapsed time to lengthen in the Placement phase of the MMC daily cycle, because for every data set that passes preliminary criteria checks, we also perform:

1. A LOCATE SVC to ensure the data set is properly cataloged.
2. An ENQ SVC to ensure the data set is not currently allocated (for example, to a long running job like CICS).
3. A random WRITE back to the MMC database to flag the data set as RECOMMENDED for MIGRATE.

Chapter 5

Tailoring CA-ASTEX Parameters

Creating Parameter Entries

Members of the ASX.CNTL data set contain sample parameter entries that you can tailor:

- ASXPARM contains a comprehensive example of parameter entries for the Performance Management components (PMC) and the Migration Manager component (MMC).
- ASXPARMP contains a comprehensive example of parameter entries related only to PMC.
- ASXPARMM contains a comprehensive example of parameter entries related only to MMC.
- ASXPARMS contains a comprehensive example of parameter entries related only to the Server component of CA-ASTEX (see "CA-ASTEX Server Parameters" in this chapter).

Direct your attention to the example that best applies to your environment for the simplest implementation. Since MMC-only parameters are ignored by the Performance Management initialization routines, and PMC-only parameters are ignored by the Migration Manager initialization routines, you do not have to worry about mixed usage.

If you wish, you can define your CA-ASTEX parameters in another data set. It can be either a physical sequential data set or a member of a partitioned data set (like SYS1.PARMLIB). You should allocate space for it on a DASD using a logical record length of 80 and a suitable block size.

Regardless of the data set you choose, you must specify your parameters according to these rules:

- Each parameter you specify must be in a section that begins with one of the following DEFINE statements:
  - DEFINE PARMS
  - DEFINE ESCON (PMC only)

- DEFINE VOLUMES
- DEFINE CGROUP (PMC only)
- DEFINE SGROUP

Each DEFINE statement must be on a line by itself. Some parameters can be specified on the same line as the statement; others can be specified below the statement.

- Parameter entries begin in column 1 and can continue through column 71.

Column	Description
1-4	These columns contain the SMF system ID of the system the parameters are intended for. If you leave these four positions blank, any parameters you supply with the specified DEFINE statement will apply to *all systems*. Any entry that contains an asterisk in the first column is treated as a comment.
5	This column must be blank.
6-71	The DEFINE statements and parameter values must begin on or after column 6. CA-ASTEX stops reading parameter values after it encounters the first blank space; it treats any subsequent data as a comment. You must specify each parameter in this format: `KEYWORD=value` You must use commas to separate parameters entered on the same line.

- Use a continuation character in column 72 to continue CGROUP, SGROUP, EXCLUDEDSN, EXCLUDEMC, DUA, and VOL parameter definitions on multiple lines. Leave column 72 blank for all other parameter definitions.

Using PMC Parameters

The following parameters apply to the Performance Management components of CA-ASTEX.

Note: You can accept the defaults for these parameters. We suggest that you review the values and defaults for *all* parameters.

- Parameters specified for the DEFINE PARMS section:

ALERTS	DSNSMFT	MLSIZE	PAGING
ALLOCDIR	DYNSGRP	MSROBJ%	RECORD
ALLOCHSM	IDB	MUSTMSR	RESETHR
ALLOCINT	IDBWRP	PMPGRP	SECURE
CACHEOPT	IMODE	PMP	SMSPARMS
CHAINJOB	IOSAMPL	PMPSYS	TREC
COMPLEXID	JOBSMFT	PMPVOL	VOLSMFT
DEFCGRP	MAYMSR	NEVMSR	VTOCRBLD
DINTVL	MLIST	NSMSMSR	

- Parameters specified for the DEFINE ESCON section:

  PATH
  SWITCH

- Parameters specified for the DEFINE VOLUMES section:

  PMPVOL
  SCOPE
  TYPE
  VOL

- Parameters specified for the DEFINE CGROUP section:

CACHEOPT	MSROBJ%	NEVMSR
CTYPE	MUSTMSR	NICACH
DUA	NAME	NSMSMSR
IMODE	PMPGRP	VOL
MAYMSR	PMPVOL	

- Parameters specified for the DEFINE SGROUP section:

ALLOCDIR	MAYMSR	PMPGRP	SCOPE
ALLOCHSM	MSROBJ%	PMPVOL	STATIC
DUA	MUSTMSR	NEVMSR	VOL
IMODE	NAME	NSMSMSR	

Using MMC Parameters

The following parameters apply to the Migration Manager component of ASTEX.

- Parameters specified for the DEFINE PARMS section:

DYNUNIT	MOVEMODE
EXCLUDEDSN	MOVETASK
EXCLUDEMC	NONWORKDAY
FAULTDEF	NOTIFY
HSMPARM (Required)	PREFIX (Required)
INCLUDEDSN	PRIMESYS (Required)
INCLUDEMC	SORTDD
KEEPSMF	SORTSPACE
MCDS	START (Required)
MICS	WINDOW
MINAGE	

- Parameters specified for the DEFINE VOLUMES section:

SCOPE=ALL|MMC

TYPE

VOL

- Parameters specified for the DEFINE SGROUP section:

  MOVEMODE

  NAME

  PRIORITY

  TARGET

  VOL

Note: In the parameter descriptions found later in this chapter, the boxes containing MMC-only parameters are shaded gray. For example:

```
MMCPARAMETER=value
```

Making Entries in the DEFINE PARMS Section

The parameters you specify in the DEFINE PARMS section define options that have an overall effect on CA-ASTEX processing.

The following rules apply to the DEFINE PARMS section:

- Your parameter data set can have *only one* DEFINE PARMS statement.
- The DEFINE PARMS statement must appear on a line by itself.
- All parameters must be specified below the statement.
- The DEFINE PARMS statement must have the following format:

```
DEFINE PARMS
```

Parameters Below the DEFINE PARMS Statement

The following parameters are listed in alphabetical order.

ALERTS Parameter

The ALERTS parameter tells CA-ASTEX whether or not it should process and send configuration and alert data to the CA-ASTEX Server component.

```
ALERTS=option
```

Valid values:

YES	Indicates that CA-ASTEX should process and send configuration and performance alert data to the CA-ASTEX Server component.
NO	Indicates that CA-ASTEX should NOT process and send configuration and performance alert data to the CA-ASTEX Server component.

Default: `ALERTS=NO`

ALLOCDIR Parameter

This parameter tells the DASD Manager whether it should optimize new allocation requests for all storage groups and HSM recall allocation requests.

```
ALLOCDIR=option
```

Valid values:

ON	Indicates that the DASD Manager should optimize allocation requests for all storage groups. This option is effective only if the Allocation intercept is initialized through the ALLOCINT parameter.
OFF	Indicates that the DASD Manager should *not* optimize allocation requests for all storage groups. If you select this option, you can use the ALLOCDIR parameter in the DEFINE SGROUP section to have the DASD Manager optimize allocation requests for individual storage groups.

Default: `ALLOCDIR=ON`

ALLOCHSM Parameter

This parameter tells the DASD Manager whether it should optimize non-directed HSM recall allocation requests for non-SMS storage groups.

```
ALLOCHSM=option
```

Valid values:

ON	Indicates that the DASD Manager should optimize non-directed HSM recall allocation requests for non-SMS storage groups. This option is effective only if the Allocation intercept is initialized through the ALLOCINT parameter, and the Allocation Director is enabled through the ALLOCDIR parameter.

OFF — Indicates that the DASD Manager should not optimize non-directed HSM recall allocation requests for non-SMS storage groups. If you select this option, you can use the ALLOCHSM parameter in the DEFINE SGROUP section to have the DASD Manager optimize recall allocation requests for individual non-SMS storage groups.

Default: `ALLOCHSM=OFF`

ALLOCINT Parameter

This parameter tells the DASD Manager whether it should initialize the Allocation intercept when CA-ASTEX is started.

```
ALLOCINT=option
```

Valid values:

YES — Indicates that the DASD Manager should initialize the Allocation intercept when CA-ASTEX is started. Once you select this option, you can activate Allocation Optimization at the global and storage group level.

NO — Indicates that the DASD Manager should not initialize the Allocation intercept when CA-ASTEX is started. If you select this option, allocation optimization will not function at any level.

Default: `ALLOCINT=YES`

CACHEOPT Parameter

This parameter tells the Cache Manager whether it should automatically optimize cache utilization for all cache controllers.

```
CACHEOPT=option
```

Valid values:

ON — Indicates that the Cache Manager should perform optimization for all cache controllers.

OFF — Indicates that Cache Manager should not perform optimization for all cache controllers. If you select this option, you can use the CACHEOPT parameter in the DEFINE CGROUP section to have the Cache Manager perform optimization for individual cache controllers.

Default: `CACHEOPT=ON`

Note: The Cache Optimizer is not activated for 3990-6 controllers enabled for record level caching (RLC II), EMC Symmetrix ICDAs, HDS 7700, HDS Upgraded 7700, HDS 7750, HDS 9900, Amdahl Spectris, IBM RAMAC Array Subsystems, IBM RAMAC Scalable Array, IBM RAMAC Virtual Array, and IBM 2105 Enterprise Storage Server.

CHAINJOB Parameter

This parameter defines the names of jobs or started tasks that perform chained I/O processing. A chained I/O results when a channel program has imbedded seek operations. The Cache Optimizer will optimize these chained I/Os only when the controller is over-utilized. DB2 and ADABAS are two examples of started tasks that perform chained I/O processing.

```
CHAINJOB=(jobname,jobname,...)
```

Valid values:

jobname,jobname,...	Specifies the name(s) of either a job(s) or a started task(s) that performs chained I/O processing. If you want to define several names that are similar, you can specify an asterisk for any unique characters.
Default:	CA-ASTEX *automatically* detects the DB2 started task name that performs chained I/O processing. That name has the following format: ****DBM1.
Example:	`CHAINJOB=(ADA*)` defines the ADABAS started task name.

COMPLEXID Parameter

This parameter is a user-defined parameter. It provides CA-ASTEX with a logical name that it can use to group together those systems that are part of the same complex.

```
COMPLEXID=xxxxxxxx
```

Valid values:

xxxxxxxx	A name consisting of at least one, but no more than eight alphanumeric characters. The name must begin with an alpha character.
Default:	_____ (8 underscores)

DEFCGRP Parameter

This parameter indicates how the storage director ID that is used in the default contention group name (DCG*xxx*) should be displayed.

```
DEFCGRP=option
```

Valid values:

DEC	Indicates that the value will be printed as a decimal value.
HEX	Indicates that the value will be printed as a hex value.
Default:	`DEFCGRP=DEC`

DINTVL Parameter

This parameter defines an "interval length" for performance data collected in detail mode.

```
DINTVL=interval length
```

Valid values:

nnnn — Specifies an interval length in minutes. This can be any value from 10 to 1440. (If you make this a number that divides evenly into 60, or is a multiple of 60, CA-ASTEX aligns intervals to stop and start on the hour.) At the beginning of an interval, CA-ASTEX starts collecting performance data and maintains it in the managed list for the number of minutes you have specified. At the interval's end, performance data collected during the interval is optionally written to your SMF data set, an Interval Database, or both. It is then purged from the managed list.

Default: `DINTVL=60`

DSNSMFT Parameter

This parameter tells CA-ASTEX whether it should write collected data set performance data to the SMF data set.

Note: You may write performance data to the Interval Database instead of or in addition to the SMF data set (see the IDB parameter). Writing to both areas involves additional space requirements and overhead.

```
DSNSMFT=record type
```

Valid values:

xxx — Specifies an SMF record type, unique to your site, for CA-ASTEX to use when writing collected data set performance data to your SMF data set at the end of each detail mode interval. You can assign the same or different record types to your DSNSMFT, JOBSMFT, and VOLSMFT parameters. Make sure the record type is not being used by another product.

Note: If you want to use CA-ASTEX to view this data online, you must also write the data to the Interval Database (see the IDB parameter), or use the CA-ASTEX Merge Utility (see the chapter "Understanding the CA-ASTEX Facilities in the CA-ASTEX *User Guide*) before viewing to convert the data to Interval Database format.

NONE — Indicates that CA-ASTEX should not write this type of record to an SMF data set. Specify a value other than NONE if you want to use SMF as input to any of the CA-ASTEX batch components.

Default: `DSNSMFT=NONE`

DYNSGRP Parameter

This parameter indicates how often CA-ASTEX dynamically verifies and updates its SMS storage group information.

The value specified should be determined by how often SMS storage group changes are made in your environment.

```
DYNSGRP=nn
```

Valid values:

nn Specifies either the number of hours, or the interval that should pass before CA-ASTEX dynamically updates its storage group information.

Default: `DYNSGRP=24`

This default indicates that CA-ASTEX should dynamically update its storage group information once every 24 hours.

DYNUNIT and the SMS DATACLAS Constructs

SMS allows you to define a DATACLAS (data class) with an attribute that allows for allocations to be multivolume. Many sites use this attribute of DATACLAS to avoid Sx37 abends.

All of the dynamically allocated data sets created by MMC would be vulnerable to these Sx37 exposures. Therefore, to avoid these operational problems, *we recommend assigning a data class that allows for multivolume allocations*.

DYNUNIT offers a similar facility if for some reason SMS is not an option. We recommend that you code DYNUNIT=(*unitname*,10) to make Sx37 a virtual impossibility. Note that all allocation parameters can be overridden by ACS routines and the constructs they assign.

DYNUNIT Parameter

This parameter lets you have greater control over dynamic allocations made by MMC, especially if MMC data sets are *not* SMS-managed. If the DYNUNIT parameter is specified, that *unitname* and *unitcount* will be used in the dynamic allocations of all Migration Manager data sets other than the database.

```
DYNUNIT=(unitname,unitcount)
```

Valid values:

unitname Specifies an esoteric unit name defined in an IOGEN (or equivalent) to group devices under a common name. The unit name can be anything that is valid for the UNIT= parameter of a DD statement in JCL. SYSALLDA is an OS/390 internal unit name created for all DASD.

unitcount Identifies the number of units to be grouped under a common name.

Specifying a count greater than one can cause multivolume data sets to be created. If you do not want multivolume data sets, specify `DYNUNIT=(unitname,1)`.

If you want to define only a number of units to be used, specify `DYNUNIT=(*NONE*, unitcount)`.

Default: None. The Migration Manager does not provide a UNIT parameter for allocation.

EXCLUDEDSN Parameter This parameter lists one or more data sets or data set prefixes for which you do not want the Migration Manager to generate data movement commands. Multiple statements are allowed.

```
EXCLUDEDSN=data set name,...
```

Valid values:

data set name,... Specifies one or more data sets you do not want moved by the Migration Manager. If you want to define several data sets that have similar names, you can use an asterisk (*) to represent any unique suffix. When you specify a list of data sets, the list must be enclosed in parentheses.

Examples:

```
EXCLUDEDSN=SYS1.*
EXCLUDEDSN=(A.B.C,SYS2.*)
EXCLUDEDSN=(SYS1.*,                          *
           A.B.C)
EXCLUDEDSN=(SYS1.*,SYS2*,ANYNODE.X.Y.DATA)
```

Note: The asterisk (*) displayed to the far right in one of the examples is a continuation character in column 72.

Even if EXCLUDEDSN is not used, data sets will still be excluded from Migration Manager control if HSM SETMIG commands are specified in the ARCCMD*xx* member of the HSM PARMLIB. The following SETMIG command examples will cause the Migration Manager to create data set exclusions:

```
SETMIG DSN(A.B.C) NOMIGRATE
SETMIG LEVEL(A) NOMIGRATE
```

EXCLUDEMC Parameter

This parameter lists one or more SMS management classes for which you do not want the Migration Manager to generate data movement commands. Multiple statements are allowed.

```
EXCLUDEMC=management class,...
```

Valid values:

management class,...	Specifies one or more management classes you do not want managed by the Migration Manager. No data sets assigned to these management classes will be moved. If you want to define several management classes that have similar names, you can use an asterisk (*) to represent any unique suffix. When you specify a list of management classes, the list must be enclosed in parentheses.

Examples:

```
EXCLUDEMC=TECHGRP
EXCLUDEMC=(TECHGRP,SDPGRP)
EXCLUDEMC=(TECHGRP,                    *
           SDPGRP)
EXCLUDEMC=(SPECIAL*,POLICY3)
```

Note: The asterisk (*) displayed to the far right in one of the examples is a continuation character in column 72.

FAULTDEF Parameter

The FAULTDEF parameter allows you to control the definition of faults.

```
FAULTDEF=(m,r)
```

A *fault* is a reference to a data set within a certain number of days after a decision to migrate. The Migration Manager tries to minimize faults by making intelligent decisions, but it can still cause faults as it learns more about your access patterns and tries to achieve your storage objectives.

The Migration Manager can cause another type of fault. A Migration Manager *recall fault* occurs when a predictive recall is performed or recommended, and then the data set is not referenced.

Valid values:

m	Specifies the number of days *after* a decision to migrate. If a reference to the data set occurs within this period, this will be considered a fault.
r	Specifies the number of days *after* a Migration Manager decision to recall. If no reference to the data set occurs within this period, this will be considered a fault.

Default: `FAULTDEF=(5,5)`

Explanation of default values:

- If a reference occurs more than 5 days after a decision to migrate, it *will not* be considered a fault; if a reference occurs 1, 2, 3, 4, or 5 days after a decision to migrate, it *will* be considered a fault.

- If a decision to recall is made by the Migration Manager, and the data set is referenced 1, 2, 3, 4, or 5 days later, it *will not* be considered a fault; otherwise, it *will* be a fault.

Recommendation: Accept this definition as a measure of efficiency. Changing this value affects other algorithms in the product. Contact Computer Associates Technical Support before changing these values.

HSMPARM Parameter

This parameter identifies the data set and member that contains the HSM parameters. If you use a single, common set of ARCCMD*xx* parameters, only one HSMPARM parameter is needed. This is ***required*** for each OS/390 image that runs HSM. If you code multiple HSMPARM parameters without specifying CPU IDs, the last one coded will be in effect.

```
HSMPARM=data set name(member)
```

Valid values:

data set name(member) Specifies the name of the data set (and member, if necessary) that contains the HSM parameters.

Example:

```
CPU1    HSMPARM=SYS1.HSM.PARMLIB(ARCCMD01)
CPU2    HSMPARM=SYS2.HSM.PARMLIB(ARCCMD02)
```

IDB Parameter

This parameter tells CA-ASTEX whether it should write collected performance data for volumes, data sets, and jobs to the Interval Database.

```
IDB=write option
```

Note: You may write performance and utilization data to the SMF data set instead of or in addition to the Interval Database (see the DSNSMFT, JOBSMFT, and VOLSMFT parameters). Writing to both areas involves additional space requirements and overhead.

Valid values:

YES	Indicates that CA-ASTEX should write performance data to the Interval Database.
NO	Indicates that CA-ASTEX should not write performance data to the Interval Database.

Default: `IDB=YES`

IDBWRP Parameter

This parameter tells CA-ASTEX what to do when the Interval Database becomes full.

```
IDBWRP=overwrite option
```

Valid values:

YES	Indicates that CA-ASTEX should continue writing data records to a full Interval Database by overwriting the oldest records in the database.

NO	Indicates that CA-ASTEX should discontinue writing records to a full Interval Database and issue this message on the operator's console:
	`ASX447I - INTERVAL DATABASE UPDATING STOPPED.`
	`NO MORE SPACE`

Default: `IDBWRP=YES`

IMODE Parameter

This parameter defines the CA-ASTEX default performance data collection mode for each volume or group of volumes.

```
IMODE=mode
```

Note: An IMODE parameter specified in the DEFINE CGROUP section or the DEFINE SGROUP section overrides any IMODE value specified in the DEFINE PARMS section.

Valid values:

DM	Specifies *detail mode* as the CA-ASTEX default performance data collection mode. As long as CA-ASTEX runs in detail mode, it continuously collects performance data for every data set and job on each volume that it is instructed to monitor. The CA-ASTEX address space code takes the data for the most active data sets and jobs accessing each volume, and stores that data in a managed list. Data from this managed list can be displayed at your terminal, or written to the Interval Database or your SMF data set.

EM — Specifies *exception mode* as the CA-ASTEX default performance data collection mode. In exception mode, CA-ASTEX collects data set and job performance data only when your response time objective is not being met. When your response time objectives are met again, CA-ASTEX stops collecting performance data for data sets and jobs.

NONE — Indicates that CA-ASTEX should not collect detailed job or data set performance data for any volume or group of volumes.

Default: `IMODE=EM`

INCLUDEDSN Parameter

This parameter lists one or more data sets or data set prefixes for which you want the Migration Manager to generate data movement commands. Multiple statements are allowed.

```
INCLUDEDSN=data set name,...
```

Valid values:

data set name,... — Specifies one or more data sets you want moved by the Migration Manager. If you want to define several data sets that have similar names, you can use an asterisk (*) to represent any unique suffix. When you specify a list of data sets, the list must be enclosed in parentheses.

Examples:

```
INCLUDEDSN=SYS1.*
INCLUDEDSN=(A.B.C,SYS2.*)
INCLUDEDSN=(SYS1.*,                          *
           A.B.C)
INCLUDEDSN=(SYS1.*,SYS2*,ANYNODE.X.Y.DATA)
```

Note: The asterisk (*) displayed to the far right in one of the examples is a continuation character in column 72.

INCLUDEMC Parameter

This parameter lists one or more SMS management classes for which you want the Migration Manager to generate data movement commands. Multiple statements are allowed.

```
INCLUDEMC=management class,...
```

Valid values:

management class,... Specifies one or more management classes you want managed by the Migration Manager. All data sets assigned to these management classes will be moved. If you want to define several management classes that have similar names, you can use an asterisk (*) to represent any unique suffix. When you specify a list of management classes, the list must be enclosed in parentheses.

Examples:

```
INCLUDEMC=TECHGRP
INCLUDEMC=(TECHGRP,SDPGRP)
INCLUDEMC=(TECHGRP,                    *
           SDPGRP)
INCLUDEMC=(SPECIAL*,POLICY3)
```

(The asterisk (*) displayed to the far right in one of the examples is a continuation character in column 72.)

IOSAMPL Parameter

This parameter tells CA-ASTEX how many I/O operations it should sample.

```
IOSAMPL=operations
```

Valid values:

nn Specifies the number of I/O operations CA-ASTEX should analyze. This can be any value from 1 to 10. For example, if you specify IOSAMPL=1, then CA-ASTEX will examine each I/O operation. If you specify IOSAMPL=2, it examines one out of every two I/O operations. If you specify IOSAMPL=10, every tenth I/O operation will be examined. The data you see, however, is adjusted to reflect the actual number of I/Os that occurred.

Default: `IOSAMPL=1`

Note: If you are using the Cache Optimizer, setting this parameter to IOSAMPL=1 ensures that the Cache Optimizer will make well-informed decisions.

JOBSMFT Parameter

This parameter tells CA-ASTEX whether it should write collected job performance data to the SMF data set.

```
JOBSMFT=record type
```

Note: You may write performance data to the Interval Database instead of or in addition to the SMF data set (see the IDB parameter). Writing to both areas involves additional space requirements and overhead.

Valid values:

xxx Specifies an SMF record type, unique to your site, for CA-ASTEX to use when writing collected job performance data to your SMF data set at the end of each detail mode interval. You can assign the same or different record types to your DSNSMFT, JOBSMFT, and VOLSMFT parameters. Make sure the record type is not being used by another product.

Note: If you want to use CA-ASTEX to view this data online, you must also write the data to the Interval Database (see the IDB parameter), or use the CA-ASTEX Merge Utility (see the chapter "Understanding the CA-ASTEX Facilities" in the CA-ASTEX *User Guide*) before viewing to convert the data to Interval Database format.

NONE Indicates that CA-ASTEX should not write this type of record to an SMF data set. Specify a value other than NONE if you want to use SMF as input to any of the CA-ASTEX batch components.

Default: `JOBSMFT=NONE`

KEEPSMF Parameter

This parameter specifies whether the Migration Manager keeps or discards the SMF records from HSM after using them.

```
KEEPSMF=option
```

Valid values:

ON | YES | Y Indicates the Migration Manager will allow the HSM record to be written to SMF.

OFF | NO | N Indicates the Migration Manager will discard the HSM record and not write it to SMF.

Default: `KEEPSMF=YES`

MAYMSR Parameter

This parameter defines your I/O millisecond response time objective for data sets defined to the SMS may-cache storage class type.

```
MAYMSR=ttt
```

Note: A MAYMSR parameter specified in the DEFINE CGROUP section or the DEFINE SGROUP section overrides any MAYMSR value specified in the DEFINE PARMS section.

Valid values:

ttt	Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.
Default:	`MAYMSR=25` This default response time objective indicates that 90% of the I/O response times should be less than or equal to 25 milliseconds each. (The MSROBJ% parameter would also have to be set to 90.)

MCDS Parameter

This parameter specifies the name of the migration control data set (MCDS) for HSM.

```
MCDS=data set name
```

Valid values:

data set name	Identifies the MCDS (a fully-qualified data set name without quotes) for HSM.
Example:	`MCDS=SYS2.DFHSM.MCDS`

MICS Parameter

If your site runs the NeuMICS DASD Space Collector once per day, this parameter can allow the Migration Manager to perform this data collection task on behalf of NeuMICS. This way, one inventory serves two different applications.

```
MICS=option
```

Valid values:

ON | YES | Y	Indicates CA-ASTEX will *not* delete the prefix.VCCDATA and prefix.VCCHSM data sets when it has finished processing them. That is, the two output files created at the START time specified in CA-ASTEX (in the ASXPARM data set) will be kept as inputs to the NeuMICS DAILY job.
OFF | NO | N	Indicates CA-ASTEX will delete the prefix.VCCDATA and prefix.VCCHSM data sets when it no longer needs them.
Default:	`MICS=NO`

MINAGE Parameter

This parameter tells the Migration Manager *not* to migrate data sets until the data sets are inactive for the specified number of days.

```
MINAGE=days
```

Valid values:

days	Identifies a whole number of days.
Default:	MINAGE=1
Recommendation:	Avoid using MINAGE because it will undermine the effectiveness and innovations of MMC, rendering it more equivalent to HSM MGMTCLAS parameters.

MLIST Parameter

This parameter defines what type of data you want CA-ASTEX to maintain in its managed lists when operating in detail mode. The CA-ASTEX managed lists contain the performance data that CA-ASTEX collects and displays. For more information about managed lists, see the CA-ASTEX *User Guide*.

```
MLIST=information type
```

Valid values:

DSN	Specifies that you want only data set information maintained in the managed list.
JOB	Indicates that you want only job information maintained in the managed list.
ALL	Indicates that you want both job and data set information included in the managed list.
Default:	`MLIST=ALL`

MLSIZE Parameter

This parameter limits the size of the managed lists CA-ASTEX uses to maintain job and data set performance measurements. (See the CA-ASTEX *User Guide* for more information about managed lists.)

```
MLSIZE=(min,max)
```

Valid values:

min Specifies the minimum number of entries the managed list must contain before CA-ASTEX can reuse any entries. This can be any number from 5 to 46, as long as it is less than or equal to *half* of the *max* value described below.

max Specifies the maximum number of entries the managed list can contain. This can be any value from 10 to 92.

Default: `MLSIZE=(15,30)`

MOVEMODE Parameter

This parameter specifies whether the Migration Manager generates commands or makes recommendations for automatic data movement.

Note: A MOVEMODE parameter specified in the DEFINE SGROUP section overrides the MOVEMODE value specified in the DEFINE PARMS section. That is, MOVEMODE at the DEFINE PARMS level establishes a global value, whereas MOVEMODE at the DEFINE SGROUP level overrides the global value for that storage group only.

```
MOVEMODE=option
```

Valid values:

ON | YES | Y Indicates the Migration Manager will generate commands to move data to other storage levels.

OFF | NO | N Indicates the Migration Manager will make recommendations to move data to other storage levels, but will do no actual data movement.

Default: `MOVEMODE=OFF`

Recommendation: Start your use of PROCMMC with `MOVEMODE=OFF` for all storage groups. Try to find at least one storage group to be enabled for `MOVEMODE=ON` as soon as possible. When you do this, make sure you review the TARGET parameter for the chosen storage group. See the section "Excessive MMC Migration Recommendations" in the chapter "Configuring and Starting CA-ASTEX" for more information.

MOVETASK Parameter

This parameter specifies whether the Migration Manager will start move request processing on certain system IDs.

```
MOVETASK=option
```

Valid values:

ON | YES | Y	Indicates the Migration Manager will start move request processing.
OFF | NO | N	Indicates the Migration Manager will *not* start move request processing.
Default:	`MOVETASK=YES`

Usage Notes:

1. Do not specify `OFF` (or `NO` or `N`) for the system ID specified in the PRIMESYS. If a system ID is not specified in columns 1 through 4, the Migration Manager will not start on the PRIMESYS. If you do not want the Movement Task to be active on the PRIMESYS, start the Migration Manager using **MOVETASK=YES** with the PRIMESYS system ID in columns 1 through 4. Then issue the **F PROCMMC,MOVETASK=OFF** (or **NO** or **N**) command on the PRIMESYS to stop move requests from being processed. (See the CA-ASTEX *User Guide* for more information about commands.)
2. **MOVETASK=OFF** is intended for special or emergency situations only. Normally, if there is work for HSM to do (a function of **MOVEMODE=ON**), you would want **MOVETASK=YES** in effect on every possible OS/390 image running HSM. Then if the need arises, you could disable MOVETASK processing as a temporary measure.

MSROBJ% Parameter

This parameter defines the percentage of time your response time objective should be met.

Note: An MSROBJ% parameter specified in the DEFINE CGROUP section or the DEFINE SGROUP section overrides any MSROBJ% value specified in the DEFINE PARMS section.

```
MSROBJ%=pp
```

Valid values:

pp	Identifies a percentage value between 50 and 99%.
Default:	`MSROBJ%=90` This default millisecond response time objective indicates that 90% of the I/O response times should be less than or equal to the values you specify on the NSMSMSR, MUSTMSR, MAYMSR, and NEVMSR parameters.

MUSTMSR Parameter

This parameter defines your I/O millisecond response time objective for data sets defined to the SMS must-cache storage class type.

Note: A MUSTMSR parameter specified in the DEFINE CGROUP section or the DEFINE SGROUP section overrides any MUSTMSR value specified in the DEFINE PARMS section.

```
MUSTMSR=ttt
```

Valid values:

ttt Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.

Default: `MUSTMSR=10`

This default response time objective indicates that 90% of the I/O response times should be less than or equal to 10 milliseconds each. (The MSROBJ% parameter would also have to be set to `90`.)

NEVMSR Parameter

This parameter defines your I/O millisecond response time objective for data sets defined to the SMS never-cache storage class type.

Note: A NEVMSR parameter specified in the DEFINE CGROUP section or the DEFINE SGROUP section overrides any NEVMSR value specified in the DEFINE PARMS section.

```
NEVMSR=ttt
```

Valid values:

ttt Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.

Default: `NEVMSR=35`

This default response time objective indicates that 90% of the I/O response times should be less than or equal to 35 milliseconds each. (The MSROBJ% parameter would also have to be set to `90`.)

NONWORKDAY Parameter

This parameter allows you to identify the non-working days of the year. In this context, a *non-working day* is defined as a day when the computing system is operating, but the usual data center users are not present. For most installations, Saturday and Sunday are the normal non-working days. In most cases, the data center will still be operating, and may actually be quite busy running backups or other maintenance work, but the usual HSM activity will not be present. It is expected that work will be running on non-working days, and in fact, Migration Manager should still be running its daily space management cycle even on these days. NONWORKDAY allows Migration Manager to make more intelligent movement decisions, as the calculation of several key items used to control movement will be enhanced to account for the data center calendar.

The Migration Manager uses the NONWORKDAY values in two different ways:

1. When the calculated next reference date (NRD) for a data set falls on a non-working day, the NRD will be incremented to the next working day. This may cause concern that the data set will become a better candidate for migration, but the next rule will account for this.
2. When all the days between the current date and the NRD are non-working days, the data set will not be considered a candidate for migration. This will prevent the Friday migration of data sets that will be used on Monday (assuming Saturday and Sunday are defined as non-working days). Note this rule applies to migration decisions only, and will not be a factor when making recall decisions. In fact, recall decisions will be more accurate because of rule 1.

```
NONWORKDAY=day1,...
```

Valid values:

day1,... Represents one or more values that describe the non-working days for the next year. You can define dates up to 366 days from the current date, using the following types of values:

- Absolute days may be coded as *ddmmmyy* (e.g., `26DEC97`). This represents the actual date of the non-working day.
- Nonspecific days may be coded as *ddmmm* (e.g., `04JUL`). This represents a non-working day that is always the same date regardless of year.
- Relative days (that is, `MON`, `TUE`, `WED`, `THU`, `FRI`, `SAT`, `SUN`) may be coded to indicate that specific days of the week are always non-working days.
- The value `NONE` may be coded to indicate that all days should be considered working days.

Default: `NONWORKDAY=(NONE)`

Example: `NONWORKDAY=(SAT,SUN,29MAY97,04JUL,25DEC,01JAN)`

NOTIFY Parameter

The Migration Manager address space is designed to stay active at all times to ensure all FSR activity is captured and to monitor the time for the start of the MMC daily cycle. If problems occur, messages are always written to the JOURNAL file (typically a SYSOUT to a HELD output class), and some messages even get written to the console in non-roll-delete mode (to intentionally attract attention).

This enhancement extends this notification to the administrator level. By specifying **NOTIFY=(*user ID,...*)**, the MMC administrator directs the Migration Manager to issue SEND commands to the TSO user ID(s) specified. The intent is to alert a responsible user of the need to investigate the JOURNAL for the abnormal condition.

Note: This is an optional extension to all the current reporting that is done for abnormal conditions; it does not replace any current logging to either JOURNAL or SYSLOG.

```
NOTIFY=(user ID,...)
```

Valid values:

user ID,... Defines the TSO user ID(s) that will receive SEND commands issued by the Migration Manager.

Example:

```
NOTIFY=(USER001,USER002)
```

NSMSMSR Parameter

This parameter defines your I/O response time objective for volumes not defined to SMS.

Note: A NSMSMSR parameter specified in the DEFINE CGROUP section or the DEFINE SGROUP section overrides any NSMSMSR value specified in the DEFINE PARMS section.

```
NSMSMSR=ttt
```

Valid values:

ttt Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.

Default:

```
NSMSMSR=30
```

This default response time objective indicates that 90% of the I/O response times should be less than or equal to 30 milliseconds each. (The MSROBJ% parameter would also have to be set to `90`.)

PAGING Parameter

This parameter specifies whether you want CA-ASTEX to measure I/O activity for page data sets.

```
PAGING=measurement option
```

Valid values:

YES Indicates that CA-ASTEX should measure page data set I/O.

NO Indicates that CA-ASTEX should not measure page data set I/O.

Default:

```
PAGING=YES
```

PMPGRP Parameter

This statement controls the flow of real time CA-ASTEX Performance Monitor Component (PMC) group-level data from the CA-ASTEX started task to the Unicenter Performance Management Predictor (Unicenter PMP) when PMP=ON.

```
PMPGRP=option
```

```
```

Valid values:

OFF	The CA-ASTEX started task does not send PMC group-level data to Unicenter PMP.
ON	If authorized, the CA-ASTEX started task attempts to send both Cache Manager and DASD Manager group-level data to Unicenter PMP over an existing interface connection.
CACHE	If Cache Manager is authorized, the CA-ASTEX started task attempts to send Cache Manager group-level data to Unicenter PMP over an existing interface connection.
DASD	If DASD Manager is authorized, the CA-ASTEX started task attempts to send DASD Manager group-level data to Unicenter PMP over an existing interface connection.
Default:	`ON`

PMP Parameter

This parameter controls the initial flow of all real time CA-ASTEX Performance Monitor Component (PMC) data from the CA-ASTEX started task to the Unicenter PMP at the time CA-ASTEX is started. When set to ON, the specific data sent is controlled by the PMPVOL=, PMPSYS= settings on the DEFINE PARMS statement. The PMPVOL=, PMPGRP= settings on the DEFINE PARMS statement may be selectively overridden for specific volumes and groups on specific DEFINE VOLUMES, DEFINE CGROUP, DEFINE SGROUP statements.

```
PMP=option
```

Valid values:

ON	The CA-ASTEX started task attempts to send PMC data to Unicenter PMP over a newly established interface connection.
OFF	The CA-ASTEX started task does not send PMC data to Unicenter PMP.
Default:	`OFF`

PMPSYS Parameter

This statement controls the flow of real time CA-ASTEX Performance Monitor Component (PMC) system-level data from the CA-ASTEX started task to the Unicenter PMP when PMP=ON.

```
PMPSYS=option
```

Valid values:

OFF	The CA-ASTEX started task does not send PMC system-level data to Unicenter PMP.
ON	If authorized, the CA-ASTEX started task attempts to send both Cache Manager and DASD Manager system-level data to Unicenter PMP over an existing interface connection.
CACHE	If Cache Manager is authorized, the CA-ASTEX started task attempts to send Cache Manager system-level data to Unicenter PMP over an existing interface connection.
DASD	If DASD Manager is authorized, the CA-ASTEX started task attempts to send DASD Manager system-level data to Unicenter PMP over an existing interface connection.

Default: `ON`

PMPVOL Parameter

This statement controls the flow of real time CA-ASTEX Performance Monitor Component (PMC) volume-level data from the CA-ASTEX started task to the Unicenter PMP when PMP=ON.

```
PMPVOL=option
```

Valid values:

OFF	The CA-ASTEX started task does not send PMC volume-level data to Unicenter PMP.
ON	If authorized, the CA-ASTEX started task attempts to send both Cache Manager and DASD Manager volume-level data to Unicenter PMP over an existing interface connection.
CACHE	If Cache Manager is authorized, the CA-ASTEX started task attempts to send Cache Manager volume-level data to Unicenter PMP over an existing interface connection.
DASD	If DASD Manager is authorized, the CA-ASTEX started task attempts to send DASD Manager volume-level data to Unicenter PMP over an existing interface connection.

Default: `OFF`

PREFIX Parameter

This parameter designates the high-level qualifiers of the data sets and the database for the Migration Manager. The qualifiers you specify here must match the prefixes you use in the ASXALLOC job (and the ASXMMGDG job).

```
PREFIX=high-level qualifiers
```

Valid values:

high-level qualifiers	Specifies the high-level qualifiers for the Migration Manager work files. If you follow the recommendation to use ASXMMGDG, the maximum length is 27 characters. Otherwise, the maximum length is 36 characters.
Default:	None.
Example:	`PREFIX=ASX.MMC`

Important! The choice you make for the value of PREFIX is critical to the successful operation of this product. See the chapter "MMC: Planning Before Customizing" before choosing this parameter value.

PRIMESYS Parameter

This parameter identifies the system ID upon which Migration Manager daily cycle processing is executed. Normally, you should specify the same SYSID as that which HSM primary daily space management executes. PRIMESYS defines which system will be activated at the START time you specify. At that time, CA-ASTEX takes an inventory of storage, updates the database, and makes its placement decision. All other SYSIDs running the Migration Manager only collect FSR data and execute the MOVETASK.

```
PRIMESYS=sysid
```

Valid values:

sysid	Specifies the ID (that is, the SMFID) of the system that initiates daily management of the Migration Manager.
Example:	`PRIMESYS=SYSA`

RECORD Parameter

This statement controls whether Unicenter PMP will record CA-ASTEX PMC performance data for complete processing.

```
RECORD=YES|NO
```

Valid values:

YES	Unicenter PMP records CA-ASTEX PMC peformance data for complete processing.
NO	Used primarily to test the CA-ASTEX interface connection to Unicenter PMP.
Default :	`YES`

RESETHR Parameter

This parameter indicates what hour CA-ASTEX should begin its all-day data collection mode (that is, reset its all-day measurements to zero).

The value specified here will also affect the amount of all-day data available to the Allocation Director when it makes its allocation decisions. If most of your allocations occur during a specific portion of the day, you may want to reset the CA-ASTEX all-day measurements after these allocations are completed. This will provide the Allocation Director with more all-day data to use in its decision-making process.

```
RESETHR=nn
```

Valid values:

nn	Specifies the hour when CA-ASTEX should reset its all-day measurements to zero. Valid values range from 0 (12 a.m.) to 23 (11 p.m.).
Default:	`RESETHR=0` This default hour indicates that CA-ASTEX will reset its all-day measurements to zero each day at midnight.

SECURE Parameter

This parameter indicates whether the security for the Command Interface is to be activated.

```
SECURE=option
```

Valid values:

YES	Indicates that security is to be activated for the Command Interface.
NO	Indicates that security will not be activated for the Command Interface.
Default:	`SECURE=NO`

SMSPARMS Parameter

This parameter indicates whether the SMS response time parameters (MUSTMSR, MAYMSR, and NEVMSR) are to be used. This allows you to specify multiple MSR objectives for your SMS volumes.

```
SMSPARMS=option
```

Valid values:

YES	Indicates that only the SMS parameters are to be used. The NSMSMSR parameter will not be accepted for SMS volumes.
NO	Indicates that only the NSMSMSR parameter will be accepted for all volumes.
Default:	`SMSPARMS=YES`

SORTDD Parameter

This parameter dynamically allocates sort work space, so that the space is only needed during the daily cycle processing time.

```
SORTDD=(prefix,nn)
```

Valid values:

prefix	Identifies the prefix of the ddnames. The ddnames used for allocations are formed by appending a two-digit number to this prefix. The maximum length is six characters.
nn	Indicates the number of DDs to be allocated for sort work space. This value must be between 1 and 99.
Default:	`SORTDD=(SORTWK,3)` causes the SORTWK01, SORTWK02, and SORTWK03 DDs to be allocated when needed.

SORTSPACE Parameter

This parameter specifies the amount of primary space to allocate for *each* allocated DD. No secondary space is allocated for the sort work DDs. (SORTSPACE is used in conjunction with SORTDD.)

```
SORTSPACE=(CYL|TRK,nn)
```

Valid values:

CYL|TRK	Specifies whether the amount of primary space to be allocated will be measured in cylinders (`CYL`) or tracks (`TRK`).
nn	Indicates the number of cylinders or tracks to be allocated for each DD.
Example:	`SORTSPACE=(CYL,50)` causes the SORTWK01, SORTWK02 and SORTWK03 DDs to be allocated with 50 cylinders each.
	Note: This parameter has no default values. If you do not specify values, the system defaults for dynamic allocation will be used. These defaults may be too small for the large amounts of data to be sorted.

START Parameter

This parameter specifies the time the Migration Manager should begin processing every day.

```
START=hh:mm
```

Valid values:

hh:mm	Identifies the hour (*hh*) and minute (*mm*) daily processing begins, based on a 24-hour day. Valid values range from 00:01 to 23:59. There is no default; you must choose a time. (The colon separating *hh* and *mm* is optional.)
Example:	`START=23:00`
Recommendation:	Specify a time for the START parameter that is about one hour earlier than the corresponding parameter in the HSM PARMLIB. That is, if `SECONDARYSPMGMTSTART(2300)` is specified for HSM, then specify `START=2200` for the CA-ASTEX Migration Manager.

TREC Parameter

This parameter defines a default destination for the CA-ASTEX performance trace records.

```
TREC=destination
```

Valid values:

n	Specifies that trace records should be formatted and written to some SYSOUT class *n*.
SYSO	Specifies that trace records should be formatted and written to SYSOUT class A.
DS	Specifies that records should be unformatted and written to a user data set that can be used for batch processing. You must write trace records to a user data set if you plan to use the CA-ASTEX Batch Cache Simulation, Workload Reporter, TTR Reorganization utility, or PDS Reorganization utility.
Default:	`TREC=SYSO`

VOLSMFT Parameter

This parameter tells CA-ASTEX whether it should write collected volume performance data to the SMF data set.

Note: You may write performance data to the Interval Database instead of or in addition to the SMF Data Set (see the IDB parameter). Writing to both areas involves additional space requirements and overhead.

```
VOLSMFT=record type
```

Valid values:

xxx	Specifies an SMF record type, unique to your site, for CA-ASTEX to use when writing collected volume performance data to your SMF data set at the end of each detail mode interval. You can assign the same or different record types to your DSNSMFT, JOBSMFT, and VOLSMFT parameters. Make sure the record type is not being used by another product. **Note:** If you want to use CA-ASTEX to view this data online, you must also write the data to the Interval Database (see the IDB parameter), or use the CA-ASTEX Merge Utility (see the chapter "Understanding the CA-ASTEX Facilities" in the CA-ASTEX *User Guide*) before viewing to convert the data to Interval Database format.
NONE	Indicates that CA-ASTEX should not write this type of record to an SMF data set. Specify a value other than NONE if you want to use SMF as input to any of the CA-ASTEX batch components.

Default: `VOLSMFT=NONE`

VTOCRBLD Parameter

This parameter controls how frequently CA-ASTEX rebuilds its extent table.

This extent table is rebuilt each time CA-ASTEX is started. If CA-ASTEX remains up for long periods of time, it is possible for some of the information in the extent table to become outdated. This can result in CA-ASTEX associating I/O to the wrong data set name.

To eliminate this problem, you can use the VTOCRBLD parameter to cause CA-ASTEX to rebuild its extent table at more frequent intervals.

Occasionally, CA-ASTEX must read the VTOC of a volume to obtain a data set name. When it does this, it adds the data set name plus all the extent information to an internal *extent* table. The next time an I/O occurs to this data set and CA-ASTEX is unable to locate the data set name, CA-ASTEX will check the extent table to see if the cylinder and track addresses for this I/O are in the table. If they are, then the data set name is found. If not, the VTOC will be read and the data set name and extent information saved in the extent table. Saving the data set name and extent information in the table reduces the number of times that CA-ASTEX must read the VTOC when a data set name is not found.

```
VTOCRBLD=nn
```

Valid values:

nn Specifies how often (in terms of numbers of hours) the extent table is rebuilt. Valid values range from 1 to 24.

If `VTOCRBLD=12` is specified, CA-ASTEX will rebuild its extent table twice each day. If `VTOCRBLD=24` is specified, CA-ASTEX will rebuild its extent table once each day.

We recommend that the value specified be set as high as possible (that is, 12 or 24), because additional overhead is associated with rebuilding the extent table.

Default: `VTOCRBLD=24`

WINDOW Parameter

This parameter specifies a whole number of hours within which the Migration Manager can begin daily processing. If started within the specified number of hours after the START time, the Migration Manager will begin its daily cycle. If the Migration Manager does not start within the specified number of hours, the daily cycle will begin the next day at the START time.

```
WINDOW=nn
```

Valid values:

nn Identifies how many hours past the START time daily processing can begin. Valid values range from 1 to 23.

Default: `WINDOW=2`

Making Entries in the DEFINE ESCON Section

The parameters you enter in the DEFINE ESCON section specify which channel paths are attached to ESCON directors (switches). If you have ESCON directors, you must include a DEFINE ESCON statement in the parameter data set.

The following rules apply to the DEFINE ESCON section:

- The DEFINE ESCON statement applies *only* to the Performance Management components of CA-ASTEX.
- Your parameter data set can have *only one* DEFINE ESCON statement.
- The DEFINE ESCON statement must appear on a line by itself.
- Each line in the DEFINE ESCON section should contain a system ID in the first four columns, unless each system has its own parameter data set. This is required since each system usually has different channel configurations.
- The DEFINE ESCON statement must have the following format:

```
DEFINE ESCON
```

- Do not specify the DEFINE ESCON statement if you have ESCON channels but no ESCON directors.

Parameter Below the DEFINE ESCON Statement

The following parameter may be specified below the DEFINE ESCON statement:

PATH Parameter

This parameter defines what channel paths (CHPIDS) have ESCON directors installed. This is similar to the entries made in the I/O configuration program deck for your ESCON directors. You must specify this parameter if you have an ESCON director.

```
PATH=xx|(xx,xx,xx,xx),SWITCH=ss
```

Valid values:

PATH=xx|(*xx,xx,xx,xx*)	Specifies a channel path ID or a list of channel paths that belong to an ESCON director.
SWITCH=*ss*	Defines the ESCON director or the associated channel paths.

Example:

```
     DEFINE ESCON
CPU1 PATH=(A0,A5,30,46,B0),SWITCH=05
CPU1 PATH=B4,SWITCH=05
CPU1 PATH=(B6,C4,C6,CA,21,26,28),SWITCH=06
```

Making Entries in the DEFINE VOLUMES Section

The parameters you enter in the DEFINE VOLUMES section specify which volumes CA-ASTEX will monitor. If you do not have a DEFINE VOLUMES statement in your parameter data set, CA-ASTEX will monitor all online volumes.

The following rules apply to the DEFINE VOLUMES section:

- Your parameter data set can have a maximum of one DEFINE VOLUME statement per component (the Migration Manager component or the Performance Manager component).
- The DEFINE VOLUMES statement must be on a line by itself.
- Each line in the DEFINE VOLUMES section must contain a system ID in the first four columns, unless the parameters on that line apply to all systems.
- The DEFINE VOLUMES statement must have the following format:

```
DEFINE VOLUMES,TYPE=monitoring option,SCOPE=component,PMPVOL=options
```

Parameters on the DEFINE VOLUMES Statement

You can specify the following parameters on the DEFINE VOLUMES statement:

PMPVOL Parameter

For volumes in this group, this statement controls the flow of real time CA-ASTEX Performance Monitor Component (PMC) volume-level data from the CA-ASTEX started task to the Unicenter PMP when PMP=ON..

```
PMPVOL=option
```

Valid values:

OFF	The CA-ASTEX started task does not send PMC volume-level data to Unicenter PMP.
ON	If authorized, the CA-ASTEX started task attempts to send both Cache Manager and DASD Manager volume-level data to Unicenter PMP over an existing interface connection.
CACHE	If Cache Manager is authorized, the CA-ASTEX started task attempts to send Cache Manager volume-level data to Unicenter PMP over an existing interface connection.
DASD	If DASD Manager is authorized, the CA-ASTEX started task attempts to send DASD Manager volume-level data to Unicenter PMP over an existing interface connection.

Default: `OFF`

SCOPE Parameter

This parameter tells CA-ASTEX if the DEFINE VOLUMES statement applies to both the Performance Management components and the Migration Manager component, the Performance Management components only, or the Migration Manager component only.

```
SCOPE=component
```

Valid values:

ALL	Indicates that the DEFINE VOLUMES definition applies to both the PMC and the MMC. This value is also valid if you are running the PMC only or the MMC only.
PMC	Indicates that the DEFINE VOLUMES definition applies to the PMC only.
MMC	Indicates that the DEFINE VOLUMES definition applies to the MMC only.

Default: `SCOPE=ALL`

TYPE Parameter

The TYPE parameter tells CA-ASTEX how to determine which online volumes to monitor.

```
TYPE=monitoring option
```

Valid values:

INCLUDE	Indicates that you want CA-ASTEX to *include* for monitoring only those volumes listed under the DEFINE VOLUMES statement.
EXCLUDE	Indicates that you want CA-ASTEX to *exclude* from monitoring all of the volumes listed under the DEFINE VOLUMES statement. Since CA-ASTEX does not monitor VM volumes, make sure that you exclude all of your VM volumes.
PMP	Volume group shares common PMPVOL=setting.

Default: `TYPE=EXCLUDE`

Parameter Below the DEFINE VOLUMES Statement

You can specify the following parameter below the DEFINE VOLUMES statement:

VOL Parameter

The VOL parameter lists the volumes you want CA-ASTEX to monitor or to exclude from monitoring. You can repeat the VOL parameter any number of times or define any number of volumes on a single VOL parameter.

```
VOL=(vvvvvv)
```

Valid values:

vvvvvv Specifies a volume serial number (volser) or list of volume serial numbers to be monitored or to be excluded from monitoring. Since CA-ASTEX does not monitor VM volumes, make sure that you exclude all of your VM volumes.

Examples:

```
DEFINE VOLUMES,TYPE=INCLUDE,SCOPE=ALL
       VOL=(VOL001,VOL002,VOL003)
       VOL=(TSO*,IMS**P,***RES)

DEFINE VOLUMES,TYPE=EXCLUDE,SCOPE=ALL
       VOL=(VOL006)
       VOL=(VM*)
```

Controlling the Scope of a Definition

If you run both the Performance Management components and the Migration Manager component, you may want to have one definition of volumes apply to one component, while another definition applies to the other component. You can do this using the SCOPE parameter, as shown in the following example:

```
DEFINE VOLUMES,TYPE=EXCLUDE,SCOPE=PMC
       VOL=(VM*)
DEFINE VOLUMES,TYPE=INCLUDE,SCOPE=MMC
       VOL=(TEST*)
```

In this example, the Performance Management components exclude from monitoring all volume serial numbers starting with VM, and the Migration Manager includes in its data collection only those volumes beginning with TEST*.

Making Entries in the DEFINE CGROUP Section

A *contention group* consists of volumes that share a common I/O path.

Note: The chapter "PMC: Planning Before Customizing" contains more information about contention groups.

CA-ASTEX defines default contention groups in both single-CPU and multiple-CPU environments. You should rely on these defaults rather than create your own contention groups.

The parameters you specify in the DEFINE CGROUP section indicate how you want CA-ASTEX to monitor the volumes in designated contention groups. If you are using default contention groups, the parameters you specify are applied to all volumes in the group.

The following rules apply to the DEFINE CGROUP section:

- The DEFINE CGROUP statement applies *only* to the Performance Management components of CA-ASTEX.
- You must create a separate DEFINE CGROUP statement for each contention group you want to define or specify parameters for. You can use a continuation character to continue a DEFINE CGROUP statement on multiple lines. See continuation character in the section "Creating Parameter Entries" for details on using a continuation character.
- Each line in the DEFINE CGROUP section must contain a system ID in the first four columns, unless the parameters on that line apply to all systems.
- To activate any changes you make in the DEFINE CGROUP section, you must:
  1. Stop and restart CA-ASTEX.
  2. Update the Trend Database, specifying FORMAT=YES. See the CA-ASTEX *User Guide* for information on updating the Trend Database.

Creating Your Own Contention Groups

Although the CA-ASTEX default contention groups are sufficient in most cases, you must define your own contention groups if:

- You are running OS/390 under VM.
- The storage director IDs for a storage director are not unique within your configuration.
- All CPUs attached to a Logical Control Unit (LCU) are not cabled to the same storage director.

- You want to specify the CTYPE parameter and have one of the following controllers:
  - HDS 7750 Array Subsystem
  - HDS 7700 Scalable Array that emulates a 3990-3
  - Upgraded HDS 7700 Scalable Array
  - Amdahl Spectris

The following shows the format of the DEFINE CGROUP section. You can specify parameters in any order.

```
DEFINE CGROUP,NAME=group name
       [,CACHEOPT=optimization option]
       [,CTYPE=controller type] [,IMODE=mode]
       [,MAYMSR=ttt] [,MSROBJ%=pp] [,MUSTMSR=ttt] [,PMPGRP=option [,PMPVOL=option]
       [,NEVMSR=ttt] [,NICACH=cache option] [,NSMSMSR=ttt]
```

Parameters on the DEFINE CGROUP Statement

The following parameters are listed in alphabetical order.

CACHEOPT Parameter

This parameter tells Cache Manager whether to optimize cache utilization for the contention group.

```
CACHEOPT=optimization option
```

Valid values:

ON	Indicates that cache optimization will be performed for all volumes within the contention group. This option is valid only if optimization is authorized, active, and designated for a cache controller.
OFF	Indicates that cache optimization will not be performed for the contention group.

Default: `CACHEOPT=OFF`

Note: The Cache Optimizer is not activated for 3990-6 controllers enabled for record level caching (RLC II), EMC Symmetrix ICDAs, HDS 7700, HDS Upgraded 7700, HDS 7750, HDS 9900, Amdahl Spectris, IBM RAMAC Array Subsystems, IBM RAMAC Scalable Array, IBM RAMAC Virtual Array, and IBM 2105 Enterprise Storage Server.

CTYPE Parameter

This optional parameter should be specified for the following storage subsystems:

- Amdahl Spectris
- HDS 7750 Array Subsystem
- HDS 7700 Scalable Array
- HDS Upgraded 7700 Scalable Array
- HDS 9900

CA-ASTEX is unable to automatically identify these storage subsystems. If specified on a DEFINE CGROUP statement, this parameter informs CA-ASTEX of the controller type and enables it to provide better statistics and more accurate problem and solution analysis for these subsystems.

Note: In order to use this parameter, you must define your own contention groups (cgroups) for these subsystems.

```
CTYPE=controller type
```

Valid values:

AMDSPEC	An Amdahl Spectris Subsystem
HDS7750	An HDS 7750 Array Subsystem
HDS7700	An HDS 7700 Scalable Array
HDSU7700	An Upgraded HDS 7700 Scalable Array
HDS9900	An HDS 9900 Type Device

IMODE Parameter

This optional parameter specifies a data collection mode (exception or detail mode). This mode applies to *all* members of the contention group unless otherwise specified below the statement line for a particular member. (Later, in this section, see another IMODE description for more about specifying the IMODE parameter below the statement line.)

Note: An IMODE parameter in this section overrides any IMODE value specified in the DEFINE PARMS section. Also, an IMODE parameter in the DEFINE SGROUP section overrides all IMODE values in the DEFINE CGROUP section.

```
IMODE=mode
```

Valid values:

DM	Indicates that CA-ASTEX is to run in *detail mode* to continuously collect performance data for the contention group.
EM	Indicates that CA-ASTEX is to run in *exception mode* to collect performance data for all members of the contention group only if the response time objective is not being met. Exception mode is useful only when you specify it for ***all*** volumes that share the same I/O path. Specifying exception mode for individual volumes will not help you find and solve path-related problems.
NONE	Indicates that CA-ASTEX should not collect any performance data for the contention group.
Default:	Value specified in the DEFINE PARMS section.
Example:	Detail mode is specified for contention group *3390S*, overriding the exception mode set in the DEFINE PARMS section:

```
DEFINE PARMS
       IMODE=EM
*
DEFINE CGROUP,NAME=3390S,IMODE=DM
*
```

(For more detailed information about the IMODE parameter, see the previous section "DYNUNIT and the SMS DATACLAS Constructs.")

MAYMSR Parameter

This parameter defines your I/O millisecond response time objective for data sets defined to the SMS may-cache storage class type.

Note: A MAYMSR parameter specified in the DEFINE CGROUP section overrides any MAYMSR value specified in the DEFINE PARMS section. Also, a MAYMSR parameter in the DEFINE SGROUP section overrides any MAYMSR value specified in the DEFINE CGROUP section.

```
MAYMSR=ttt
```

Valid values:

ttt	Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.
Default:	`MAYMSR=25` This default response time objective indicates that 90% of the I/O response times should be less than or equal to 25 milliseconds each. (The MSROBJ% parameter would also have to be set to `90`.)

MSROBJ% Parameter

This parameter defines the percentage of time your response time objective should be met.

Note: A MSROBJ% parameter specified in the DEFINE CGROUP section overrides any MSROBJ% value specified in the DEFINE PARMS section. Also, a MSROBJ% parameter in the DEFINE SGROUP section overrides any MSROBJ% value specified in the DEFINE CGROUP section.

```
MSROBJ%=pp
```

Valid values:

pp Identifies a percentage value between 50 and 99%.

Default: `MSROBJ%=90`

This default millisecond response time objective indicates that 90% of the I/O response times should be less than or equal to the values you specify on the NSMSMSR, MUSTMSR, MAYMSR, and NEVMSR parameters.

MUSTMSR Parameter

This parameter defines your I/O millisecond response time objective for data sets defined to the SMS must-cache storage class type.

Note: A MUSTMSR parameter specified in the DEFINE CGROUP section overrides any MUSTMSR value specified in the DEFINE PARMS section. Also, a MUSTMSR parameter in the DEFINE SGROUP section overrides any MUSTMSR value specified in the DEFINE CGROUP section.

```
MUSTMSR=ttt
```

Valid values:

ttt Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.

Default: `MUSTMSR=10`

This default response time objective indicates that 90% of the I/O response times should be less than or equal to 10 milliseconds each. (The MSROBJ% parameter would also have to be set to `90`.)

NAME Parameter

This parameter designates a contention group that includes the volumes identified under the statement.

```
NAME=group name
```

Valid values:

group name — Specifies a *contention group name* (one to eight alphanumeric characters). Each contention group name must be unique. You can find the names of CA-ASTEX-assigned contention groups on any Group/Volume screen or report after CA-ASTEX has been running for a few minutes. You may use the assigned contention group name on this parameter or substitute a name that conforms with the naming conventions at your installation.

Example: `DEFINE CGROUP NAME=DCG001`

Default: `ON`

NEVMSR Parameter

This parameter defines your I/O millisecond response time objective for data sets defined to the SMS never-cache storage class type.

Note: A NEVMSR parameter specified in the DEFINE CGROUP section overrides any NEVMSR value specified in the DEFINE PARMS section. Also, a NEVMSR parameter in the DEFINE SGROUP section overrides any NEVMSR value specified in the DEFINE CGROUP section.

```
NEVMSR=ttt
```

Valid values:

ttt — Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.

Default: `NEVMSR=35`

This default response time objective indicates that 90% of the I/O response times should be less than or equal to 35 milliseconds each. (The MSROBJ% parameter would also have to be set to `90`.)

NICACH Parameter

This parameter indicates whether the volumes in the contention group are on a controller that is not fully 3990-3 compatible. What you specify on this parameter applies to *all* volumes in the contention group unless otherwise specified for a particular volume. (Later in this section see more information about specifying the NICACH Parameter below the statement line.)

```
NICACH=cache option
```

Valid values:

YES — Indicates that the volumes in the contention group are being cached. You must specify `YES` if this is a controller that is not fully 3990-3 compatible.

NO — Indicates that the volumes in the contention group are not being cached.

DFW — Indicates that the volumes in the contention group are cached and have DASD Fast Write (DFW) turned on. (Make sure that the controller supports DFW.)

Default: `NICACH=NO`

NSMSMSR Parameter

This parameter defines your I/O response time objective for volumes not defined to SMS.

Note: An NSMSMSR parameter in the DEFINE CGROUP section overrides any NSMSMSR value specified in the DEFINE PARMS section. Also, an NSMSMSR parameter in the DEFINE SGROUP section overrides any NSMSMSR value specified in the DEFINE CGROUP section.

```
NSMSMSR=ttt
```

Valid values:

ttt — Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.

Default: `NSMSMSR=30`

This default response time objective indicates that 90% of the I/O response times should be less than or equal to 30 milliseconds each. (The MSROBJ% parameter would also have to be set to `90`.)

Example: The contention group *3390S* has a response time objective of 30 milliseconds, overriding the 25-millisecond response time objective set in the DEFINE PARMS section for all non-SMS volumes in the contention group.

```
DEFINE PARMS
       MSROBJ%=90
       NSMSMSR=25
*
DEFINE CGROUP,NAME=3390S,NSMSMSR=30
*
```

PMPGRP Parameter

For this CGROUP, this statement controls the flow of real time CA-ASTEX Performance Monitor Component (PMC) group-level data from the CA-ASTEX started task to the Unicenter PMP when PMP=ON.

```
PMPGRP=option
```

Valid values:

OFF — The CA-ASTEX started task does not send PMC group-level data to Unicenter PMP.

ON — If authorized, the CA-ASTEX started task attempts to send both Cache Manager and DASD Manager group-level data to Unicenter PMP over an existing interface connection.

CACHE	If Cache Manager is authorized, the CA-ASTEX started task attempts to send Cache Manager group-level data to Unicenter PMP over an existing interface connection.
DASD	If DASD Manager is authorized, the CA-ASTEX started task attempts to send DASD Manager group-level data to Unicenter PMP over an existing interface connection.

Default: `ON`

PMPVOL Parameter

For volumes in this CGROUP, this statement controls the flow of real time CA-ASTEX Performance Monitor Component (PMC) volume-level data from the CA-ASTEX started task to the Unicenter PMP when PMP=ON.

```
PMPVOL=option
```

Valid values:

OFF	The CA-ASTEX started task does not send PMC volume-level data to Unicenter PMP.
ON	If authorized, the CA-ASTEX started task attempts to send both Cache Manager and DASD Manager volume-level data to Unicenter PMP over an existing interface connection.
CACHE	If Cache Manager is authorized, the CA-ASTEX started task attempts to send Cache Manager volume-level data to Unicenter PMP over an existing interface connection.
DASD	If DASD Manager is authorized, the CA-ASTEX started task attempts to send DASD Manager volume-level data to Unicenter PMP over an existing interface connection.

Default: `OFF`

Parameters Below the DEFINE CGROUP Statement

The following parameters are listed in alphabetical order.

DUA Parameter

This parameter lists members of a contention group by their *device unit addresses*. You can repeat the DUA parameter as many times as necessary and specify any number of unit addresses on it. For default contention groups, it is only necessary to list one of the volumes in the group.

Note: You *cannot* specify both DUA and VOL parameters under a single DEFINE CGROUP statement; use one or the other. (See later in this section more information about the VOL parameter.)

```
DUA=(nnnn)
```

Valid values:

nnnn Specifies a device unit address. If you want to define several device unit addresses with similar names, you can substitute an asterisk in place of specific characters.

Do not assign a particular DUA (or corresponding volser) to more than one contention group. CA-ASTEX will not analyze contention groups that have common members.

Do not specify optional channel path addresses on DUA parameters. Specify primary device unit addresses only, and be sure to define DASD device addresses. CA-ASTEX does not manage MSS virtual unit addresses (VUAs).

Example 1:

```
DEFINE CGROUP,NAME=3390S
       DUA=(13*,14*)
       DUA=(1550)
```

Example 2:

```
DEFINE CGROUP,NAME=3390S,NICACH=DFW
       DUA=(961,962,963,964),NICACH=YES
       DUA=(965,966,967,968),NICACH=YES
       DUA=(969,96A,96B,96C)
       DUA=(96D,96E,96F),NICACH=NO
```

IMODE Parameter

This optional parameter specifies a data collection mode (EM, DM, or NONE) for one or more *particular* members of the contention group.

Note: An IMODE parameter that is specified *under* the DEFINE CGROUP statement overrides any IMODE value specified *on* the DEFINE CGROUP statement or in the DEFINE PARMS section. Also, an IMODE parameter in a DEFINE SGROUP section overrides all IMODE values in the DEFINE CGROUP section.

```
IMODE=mode
```

Valid values:

DM Indicates that CA-ASTEX is to run in *detail mode* to continuously collect performance data for the storage group.

EM Indicates that CA-ASTEX is to run in *exception mode* to collect performance data for all members of the storage group only if the response time objective is not being met. Exception mode is useful only when you specify it for *all* volumes that share the same I/O path. Specifying exception mode for individual volumes will not help you find and solve path-related problems.

NONE Indicates that CA-ASTEX should not collect any performance data for the storage group.

Example:

Detail mode is specified for contention group *3390S*, overriding the exception mode set in the DEFINE PARMS section:

```
DEFINE PARMS
       IMODE=EM
*
DEFINE CGROUP,NAME=3390S,IMODE=DM
```

(For more information about the IMODE parameter, see the section "PMPGRP Parameter For this CGROUP, this statement controls the flow of real time CA-ASTEX Performance Monitor Component (PMC) group-level data from the CA-ASTEX started task to the Unicenter PMP when PMP=ON.

```
PMPGRP=option
```

Valid values:

OFF	The CA-ASTEX started task does not send PMC group-level data to Unicenter PMP.
ON	If authorized, the CA-ASTEX started task attempts to send both Cache Manager and DASD Manager group-level data to Unicenter PMP over an existing interface connection.
CACHE	If Cache Manager is authorized, the CA-ASTEX started task attempts to send Cache Manager group-level data to Unicenter PMP over an existing interface connection.
DASD	If DASD Manager is authorized, the CA-ASTEX started task attempts to send DASD Manager group-level data to Unicenter PMP over an existing interface connection.

Default: `ON`

PMPVOL Parameter

For volumes in this CGROUP, this statement controls the flow of real time CA-ASTEX Performance Monitor Component (PMC) volume-level data from the CA-ASTEX started task to the Unicenter PMP when PMP=ON.

```
PMPVOL=option
```

Valid values:

OFF	The CA-ASTEX started task does not send PMC volume-level data to Unicenter PMP.
ON	If authorized, the CA-ASTEX started task attempts to send both Cache Manager and DASD Manager volume-level data to Unicenter PMP over an existing interface connection.
CACHE	If Cache Manager is authorized, the CA-ASTEX started task attempts to send Cache Manager volume-level data to Unicenter

PMP over an existing interface connection.

DASD — If DASD Manager is authorized, the CA-ASTEX started task attempts to send DASD Manager volume-level data to Unicenter PMP over an existing interface connection.

Default: `OFF`

Parameters Below the DEFINE CGROUP Statement ."

MAYMSR Parameter

This parameter defines your I/O millisecond response time objective for data sets defined to the SMS may-cache storage class type.

Note: A MAYMSR parameter that is specified *under* the DEFINE CGROUP statement overrides any MAYMSR value *on* the DEFINE CGROUP statement or in the DEFINE PARMS section. Also, a MAYMSR parameter in a DEFINE SGROUP section overrides all MAYMSR values in the DEFINE CGROUP section.

```
MAYMSR=ttt
```

Valid values:

ttt — Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.

Default: `MAYMSR=25`

MSROBJ% Parameter

This parameter defines the percentage of time your response time objective should be met.

Note: A MSROBJ% parameter that is specified *under* the DEFINE CGROUP statement overrides any MSROBJ% value *on* the DEFINE CGROUP statement or in the DEFINE PARMS section. Also, a MSROBJ% parameter in a DEFINE SGROUP section overrides all MSROBJ% values in the DEFINE CGROUP section.

```
MSROBJ%=pp
```

Valid values:

pp — Identifies a percentage value between 50 and 99%.

Default: `MSROBJ%=90`

This default millisecond response time objective indicates that 90% of the I/O response times should be less than or equal to the values you specify on the NSMSMSR, MUSTMSR, MAYMSR, and NEVMSR parameters.

MUSTMSR Parameter

This parameter defines your I/O millisecond response time objective for data sets defined to the SMS must-cache storage class type.

Note: A MUSTMSR parameter that is specified *under* the DEFINE CGROUP statement overrides any MUSTMSR value *on* the DEFINE CGROUP statement or in the DEFINE PARMS section. Also, a MUSTMSR parameter in a DEFINE SGROUP section overrides all MUSTMSR values in the DEFINE CGROUP section.

```
MUSTMSR=ttt
```

Valid values:

ttt Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.

Default: `MUSTMSR=10`

NEVMSR Parameter

This parameter defines your I/O millisecond response time objective for data sets defined to the SMS never-cache storage class type.

Note: A NEVMSR parameter that is specified *under* the DEFINE CGROUP statement overrides any NEVMSR value *on* the DEFINE CGROUP statement or in the DEFINE PARMS section. Also, a NEVMSR parameter in a DEFINE SGROUP section overrides all NEVMSR values in the DEFINE CGROUP section.

```
NEVMSR=ttt
```

Valid values:

ttt Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.

Default: `NEVMSR=35`

Example: Volume *SYSRR1* has a response time objective of 35 milliseconds, overriding the 25-millisecond response time objective set in the DEFINE PARMS section:

```
DEFINE PARMS
       MSROBJ%=90
       NEVMSR=25
DEFINE CGROUP,NAME=3465T
       VOL=(VOL003,VOL004)
       VOL=(SYSRR1),NEVMSR=35
```

NICACH Parameter

This parameter indicates whether a *particular* volume in the contention group is on a controller that is not fully 3990-3 compatible. A NICACH parameter that is specified for a volume *under* the DEFINE CGROUP statement overrides any NICACH value specified for a contention group *on* the DEFINE CGROUP statement.

```
NICACH=cache option
```

Valid values:

YES — Indicates that the volumes in the contention group are being cached. You must specify `YES` if this is a cache controller that is not fully 3990-3 compatible and the volumes are turned on to cache.

NO — Indicates that the volumes in the contention group are not being cached.

DFW — Indicates that the volumes in the contention group have DASD Fast Write (DFW) turned on. (Make sure the controller supports DFW.)

Default: `NICACH=NO`

NSMSMSR Parameter

This parameter defines your I/O response time objective for volumes not defined to SMS.

Note: A NSMSMSR parameter that is specified *under* the DEFINE CGROUP statement overrides any NSMSMSR value *on* the DEFINE CGROUP statement or in the DEFINE PARMS section. Also, a NSMSMSR parameter in a DEFINE SGROUP section overrides all NSMSMSR values in the DEFINE CGROUP section.

```
NSMSMSR=ttt
```

Valid values:

ttt — Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.

Default: `NSMSMSR=30`

This default response time objective indicates that 90% of the I/O response times should be less than or equal to 30 milliseconds each. (The MSROBJ% parameter would also have to be set to `90`.)

Example: The contention group *3390S* has a response time objective of 30 milliseconds, overriding the 25-millisecond response time objective set in the DEFINE PARMS section for all non-SMS volumes in the contention group.

```
DEFINE PARMS
       MSROBJ%=90
       NSMSMSR=25
*
DEFINE CGROUP,NAME=3390S
       VOL=(VOL001,VOL002,VOL003),NSMSMSR=30
*
```

VOL Parameter

This parameter lists members of a contention group by their *volume serial numbers* (volsers). You can repeat the VOL parameter as many times as necessary and define any number of volumes on a single VOL parameter. For default contention groups, it is only necessary to list one of the volumes in the group.

Note: You *cannot* specify both VOL and DUA parameters under a single DEFINE CGROUP statement; use one or the other. (See earlier in this section more information about the DUA parameter.)

```
VOL=(vvvvvv)
```

Valid values:

vvvvvv
: Specifies a volume serial number. If you want to define several volumes with similar volsers, you can substitute an asterisk in place of specific characters.

  Do not assign a particular volser (or corresponding DUA) to more than one contention group. CA-ASTEX will not analyze contention groups that have common members.

Example:

```
DEFINE CGROUP,NAME=3390S
       VOL=(SYS*)
       VOL=(VOL001,VOL002,VOL003)
```

Making Entries in the DEFINE SGROUP Section

A storage group consists of volumes that contain similar data or provide service to the same applications. Storage groups are necessary for the Storage Performance Expert; it recommends data set moves only between volumes in the same storage group.

If you are using SMS, CA-ASTEX defines the correct SMS storage groups for you. You should rely on these default storage group assignments rather than making your own. You can modify a default storage group by specifying your changes in a DEFINE SGROUP section. If you are not using SMS, or you are using an SMS equivalent, you must create your own storage groups. The CA-ASTEX storage group definitions must be the same as the volume pooling groups defined in the SMS-equivalent product.

The parameters you specify in the DEFINE SGROUP section indicate how you want CA-ASTEX to monitor the volumes in storage groups.

The following rules apply to the DEFINE SGROUP section:

- You must create a separate DEFINE SGROUP statement for each storage group you want to define. You can use a continuation character to continue a DEFINE SGROUP statement on multiple lines. See the section "Creating Parameter Entries" for details on using a continuation character.
- You must define at least two volumes to each DEFINE SGROUP statement if you want the Storage Performance Expert to offer a solution.

  Note: This rule applies only to the Performance Management components.
- Each line in the DEFINE SGROUP section must contain a system ID in the first four columns, unless the parameters on that line apply to all systems.
- To activate any changes you make in the DEFINE SGROUP section, you must:

  1. Stop and restart CA-ASTEX.
  2. Update the Trend Database, specifying FORMAT=YES. (See the chapter "Understanding the CA-ASTEX Facilities" in the CA-ASTEX *User Guide* for information on updating the Trend Database.)

     Note: This step applies only to the Performance Management components.

The following shows the format of the DEFINE SGROUP section. You can specify parameters in any order.

```
DEFINE SGROUP,NAME=group name [,IMODE=mode ]
           [ ,ALLOCDIR=option] [ ,ALLOCHSM=option ] [ ,MAYMSR=ttt ]

           [ ,MSROBJ%=pp ] [ ,MUSTMSR=ttt] [ ,PMPGRP=option][ ,PMPVOL=option]

           [ ,NEVMSR=ttt] [ ,NSMSMSR=ttt]

           [ ,STATIC=option ]

           [ ,TARGET=pp  ] [ ,PRIORITY=nn ]

           [ ,MOVEMODE=option ]
```

Special Storage Groups Created by the Migration Manager

The Migration Manager may create up to four additional storage groups to account for volumes within your configuration:

<ML1> Represents Migration Level 1 volumes found on your system and defined in the ARCCMD*xx* HSM PARMLIB member as Level 1 migration volumes.

<ML2> Represents those data sets residing on Migration Level 2 volumes that are known to the Migration Manager. Because these represent only a subset of all your ML2 data sets (the more active ones), this storage group will appear only on certain displays.

<NONSMS> Represents volumes found on your system that are not managed by SMS but are defined in the ARCCMD*xx* HSM PARMLIB member as primary volumes.

<NONHSM> Represents volumes found on your system that are not managed by SMS and are not defined in the ARCCMD*xx* HSM PARMLIB member.

The Migration Manager allows the target threshold percentage of occupancy and the MOVEMODE to be overridden for the default storage group <NONSMS>. A priority for processing move requests can also be set using the PRIORITY parameter. For the <ML1>, <NONSMS>, and <NONHSM> storage groups, a VOL or DUA parameter may not be coded.

You can override the target threshold percentage of occupancy for the <NONHSM> special group. However, PRIORITY and MOVEMODE will always be ignored for <NONHSM>, which can never be managed. You can manage volumes from <NONHSM> by using either the VOL or the DUA parameter to assign them to other groups.

You may not specify MOVEMODE for the <ML1> storage group. It will always default to MOVEMODE=OFF when MOVEMODE=ON is not active for any other storage group. As soon as MOVEMODE=ON is coded for *any* storage group, MOVEMODE=ON will also apply to the <ML1> group. Processing priority is also ignored when coded for this group.

Parameters on the DEFINE SGROUP Statement

The following parameters are listed in alphabetical order.

ALLOCDIR Parameter

This parameter tells the DASD Manager whether the Allocation Director should optimize new allocation requests for this storage group. To see that non-SMS volumes get the maximum benefit from the Allocation Director, make sure that you define only one SGROUP for each generic and esoteric device grouping.

```
ALLOCDIR=option
```

Valid values:

ON	Indicates that DASD Manager should optimize new allocation requests for all volumes in this storage group. This option is effective only if the Allocation intercept is initialized through the ALLOCINT parameter in the DEFINE PARMS section.
OFF	Indicates that the DASD Manager should *not* optimize new allocation requests for all volumes in this storage group.

Default: `ALLOCDIR=ON`

ALLOCHSM Parameter

This parameter tells the DASD Manager whether it should optimize HSM recall allocation requests for non-SMS volumes or volumes that are being managed as non-SMS.

```
ALLOCHSM=option
```

Valid values:

ON	Indicates that the DASD Manager should optimize HSM recall allocation requests for all volumes in this non-SMS storage group. This option is effective only if the Allocation intercept is initialized through the ALLOCINT parameter in the DEFINE PARMS section and the Allocation Director is enabled through the DEFINE PARMS section or on the DEFINE SGROUP parameter using the ALLOCDIR parameter.
OFF	Indicates that the DASD Manager should *not* optimize HSM recall allocation requests for volumes in this non-SMS storage group.

Default: `ALLOCHSM=OFF`

IMODE Parameter

This optional parameter specifies a data collection mode (exception or detail mode). This mode applies to *all* members of the storage group unless otherwise specified below the statement line for a particular member. (See the section "PMPGRP Parameter For this CGROUP, this statement controls the flow of real time CA-ASTEX Performance Monitor Component (PMC) group-level data from the CA-ASTEX started task to the Unicenter PMP when PMP=ON.

```
PMPGRP=option
```

Valid values:

OFF	The CA-ASTEX started task does not send PMC group-level data to Unicenter PMP.
ON	If authorized, the CA-ASTEX started task attempts to send both Cache Manager and DASD Manager group-level data to Unicenter PMP over an existing interface connection.
CACHE	If Cache Manager is authorized, the CA-ASTEX started task attempts to send Cache Manager group-level data to Unicenter PMP over an existing interface connection.
DASD	If DASD Manager is authorized, the CA-ASTEX started task attempts to send DASD Manager group-level data to Unicenter PMP over an existing interface connection.

Default: `ON`

PMPVOL Parameter

For volumes in this CGROUP, this statement controls the flow of real time CA-ASTEX Performance Monitor Component (PMC) volume-level data from the CA-ASTEX started task to the Unicenter PMP when PMP=ON.

```
PMPVOL=option
```

Valid values:

OFF	The CA-ASTEX started task does not send PMC volume-level data to Unicenter PMP.
ON	If authorized, the CA-ASTEX started task attempts to send both Cache Manager and DASD Manager volume-level data to Unicenter PMP over an existing interface connection.
CACHE	If Cache Manager is authorized, the CA-ASTEX started task attempts to send Cache Manager volume-level data to Unicenter PMP over an existing interface connection.
DASD	If DASD Manager is authorized, the CA-ASTEX started task attempts to send DASD Manager volume-level data to Unicenter PMP over an existing interface connection.

Default: `OFF`

Parameters Below the DEFINE CGROUP Statement" for more about specifying the IMODE parameter below the statement line.) If you do not specify a data collection mode for storage group members, CA-ASTEX will use whatever IMODE values are assigned to them elsewhere in a DEFINE CGROUP or DEFINE PARMS section.

Note: An IMODE parameter in this section overrides any IMODE values specified in the DEFINE CGROUP and DEFINE PARMS sections.

Important! *Be careful when specifying an IMODE value for members of a storage group. It might conflict with the IMODE value of a contention group, where some of the storage group volumes also may be members. All members of a contention group should have the same IMODE value.*

```
IMODE=mode
```

Valid values:

DM	Indicates that CA-ASTEX is to run in *detail mode* to continuously collect performance data for the storage group.
EM	Indicates that CA-ASTEX is to run in *exception mode* to collect performance data for all members of the storage group only if the response time objective is not being met. Exception mode is useful only when you specify it for *all* volumes that share the same I/O path. Specifying exception mode for individual volumes will not help you find and solve path-related problems.
NONE	Indicates that CA-ASTEX should not collect any performance data for the storage group.

Default: Value specified in the DEFINE PARMS section.

For more detailed information about the IMODE parameter, see the section "Parameters on the DEFINE CGROUP Statement."

MAYMSR Parameter

This parameter defines your I/O millisecond response time objective for data sets defined to the SMS may-cache storage class type.

Note: A MAYMSR parameter in this section overrides any MAYMSR values specified in the DEFINE CGROUP and DEFINE PARMS sections.

```
MAYMSR=ttt
```

Valid values:

ttt	Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.
Default:	`MAYMSR=25`

MOVEMODE Parameter

This parameter specifies whether the Migration Manager generates commands or makes recommendations for data movement.

See Excessive MMC Migration Recommendations in Chapter 6 for more consideration about the MOVEMODE and TARGET parameters. Also, refer to the description of the TARGET Parameter in the section "Parameters on the DEFINE SGROUP Statement."

Note: A MOVEMODE parameter in this section overrides the MOVEMODE value specified in the DEFINE PARMS section. That is, MOVEMODE at the DEFINE PARMS level establishes a global value, whereas MOVEMODE at the DEFINE SGROUP level overrides the global value for that storage group only.

```
MOVEMODE=option
```

Valid values:

ON | YES | Y	Indicates the Migration Manager will generate commands to move data to other storage levels.
OFF | NO | N	Indicates the Migration Manager will make recommendations to move data to other storage levels, but will do no actual data movement.
Default:	`MOVEMODE=OFF`

MSROBJ% Parameter

This parameter defines the percentage of time your response time objective should be met.

Note: A MSROBJ% parameter in this section overrides any MSROBJ% values specified in the DEFINE CGROUP and DEFINE PARMS sections.

```
MSROBJ%=pp
```

Valid values:

pp Identifies a percentage value between 50 and 99%.

Default: `MSROBJ%=90`

This default millisecond response time objective indicates that 90% of the I/O response times should be less than or equal to the values you specify on the NSMSMSR, MUSTMSR, MAYMSR, and NEVMSR parameters.

MUSTMSR Parameter

This parameter defines your I/O millisecond response time objective for data sets defined to the SMS must-cache storage class type.

Note: A MUSTMSR parameter in this section overrides any MUSTMSR values specified in the DEFINE CGROUP and DEFINE PARMS sections.

```
MUSTMSR=ttt
```

Valid values:

ttt Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.

Default: `MUSTMSR=10`

NAME Parameter

This parameter designates a storage group that includes the volumes identified under the statement.

```
NAME=group name
```

Valid values:

group name Specifies a storage group name (one to eight alphanumeric characters). Each storage group name must be unique.

NEVMSR Parameter

This parameter defines your I/O millisecond response time objective for data sets defined to the SMS never-cache storage class type.

Note: A NEVMSR parameter in this section overrides any NEVMSR values specified in the DEFINE CGROUP and DEFINE PARMS sections.

```
NEVMSR=ttt
```

Valid values:

ttt Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.

Default: `NEVMSR=35`

NSMSMSR Parameter

This parameter defines your I/O response time objective for volumes not defined to SMS.

Note: An NSMSMSR parameter in this section overrides any NSMSMSR values specified in the DEFINE CGROUP and DEFINE PARMS sections.

```
NSMSMSR=ttt
```

Valid values:

ttt	Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.
Default:	`NSMSMSR=30` This default response time objective indicates that 90% of the I/O response times should be less than or equal to 30 milliseconds each. (The MSROBJ% parameter would also have to be set to `90`.)

PMPGRP Parameter

For this SGROUP, this statement controls the flow of real time CA-ASTEX Performance Monitor Component (PMC) group-level data from the CA-ASTEX started task to the Unicenter PMP when PMP=ON.

```
PMPGRP=option
```

Valid values:

OFF	The CA-ASTEX started task does not send PMC group-level data to Unicenter PMP.
ON	If authorized, the CA-ASTEX started task attempts to send both Cache Manager and DASD Manager group-level data to Unicenter PMP over an existing interface connection.
CACHE	If Cache Manager is authorized, the CA-ASTEX started task attempts to send Cache Manager group-level data to Unicenter PMP over an existing interface connection.
DASD	If DASD Manager is authorized, the CA-ASTEX started task attempts to send DASD Manager group-level data to Unicenter PMP over an existing interface connection.
Default:	`ON`

PMPVOL Parameter

For volumes in this SGROUP, this statement controls the flow of real time CA-ASTEX Performance Monitor Component (PMC) volume-level data from the CA-ASTEX started task to the Unicenter PMP when PMP=ON.

```
PMPVOL=option
```

Valid values:

OFF	The CA-ASTEX started task does not send PMC volume-level data to Unicenter PMP.
ON	If authorized, the CA-ASTEX started task attempts to send both Cache Manager and DASD Manager volume-level data to Unicenter PMP over an existing interface connection.
CACHE	If Cache Manager is authorized, the CA-ASTEX started task attempts to send Cache Manager volume-level data to Unicenter PMP over an existing interface connection.
DASD	If DASD Manager is authorized, the CA-ASTEX started task attempts to send DASD Manager volume-level data to Unicenter PMP over an existing interface connection.
Default:	`OFF`

PRIORITY Parameter

This parameter specifies the order of data movement processing in the final stage of the daily processing cycle for the Migration Manager. Because it may take HSM several hours to do the data movement (workload) generated by the Migration Manager, this parameter allows you to prioritize that workload. If it is more important that some storage groups achieve their free space targets before the start of a normal business day, then you can rank that relative importance with this parameter. Storage groups that have equal priority are processed in ascending order according to their names.

`PRIORITY=nn`

Valid values:

nn	Indicates the rank in which a storage group qualifies for movement. Valid values range from 00 (the lowest priority) to 99 (the highest priority).
Default:	`PRIORITY=00`

STATIC Parameter

CA-ASTEX can dynamically recognize storage group changes that occur for SMS-managed volumes. This parameter designates whether or not dynamic storage group processing should occur for the storage group coded on the DEFINE SGROUP statement.

`STATIC=option`

Valid values:

YES	Indicates that dynamic storage group processing should occur for the storage group coded on the DEFINE SGROUP statement.
NO	Indicates that dynamic storage group processing should not occur for the storage group coded on the DEFINE SGROUP statement. When NO is specified, CA-ASTEX does not allow dynamic storage group processing to override the storage group names that you have

coded in your CA-ASTEX parms.

Default: `YES`

Example:
```
DEFINE SGROUP,NAME=TMPGRP,STATIC=YES
```

TARGET Parameter

This parameter defines the allocation threshold for the storage group.

`TARGET=ppp`

Valid values:

ppp Specifies the target threshold as a percentage of the total capacity for the storage group. The valid range is 0 to 100. Coding a trailing percent sign (%) is optional.

Default: Expressed in the ISMF or HSM definition.

Usage (1):

You might want to use this parameter as a simulation tool to see how changing the target threshold of a storage group affects CA-ASTEX (operating with `MOVEMODE=OFF`), but, without affecting the current operational thresholds in use by SMS and HSM. Because CA-ASTEX does not actually move the data, the goal of achieving a given occupancy level could be simulated by adding the TARGET parameter to experiment.

Usage (2):

Another use of TARGET pertains to the case of a non-SMS pool of volumes being defined to form a storage group. In this case, there is no ISMF definition available for CA-ASTEX to automatically find, as it does for SMS storage groups. Therefore, if you have defined non-SMS pools in this way, you *must*, in turn, specify `TARGET=`*ppp* to provide the target occupancy for the storage group.

Important! *If MOVEMODE=ON is in effect globally or at the storage group level, any value specified for TARGET will have a real, not simulated, effect.*

Parameters Below the DEFINE SGROUP Statement

The following parameters are listed in alphabetical order.

DUA Parameter

This parameter lists members of a storage group by their *device unit addresses*. You can repeat the DUA parameter as many times as necessary and specify any number of unit addresses on it.

Note: You *cannot* specify both DUA and VOL parameters (see below) under a single DEFINE SGROUP statement; use one or the other.

```
DUA=(nnnn)
```

Valid values:

nnnn Specifies a device unit address. If you want to define several device unit addresses with similar names, you can substitute an asterisk in place of specific characters.

Do not assign a specific DUA (or corresponding volser) to more than one storage group. CA-ASTEX will not analyze storage groups that have members in common.

Do not specify optional channel path addresses on DUA parameters. Specify primary device unit addresses only, and be sure to define DASD device addresses. CA-ASTEX does not manage MSS virtual unit addresses (VUAs).

Example:

```
DEFINE SGROUP,NAME=APPLA
       DUA=(15*,16*)
       DUA=(1480)
```

IMODE Parameter

This optional parameter specifies a data collection mode (exception or detail mode) for one or more *particular* members of the storage group.

Note: An IMODE parameter that is specified *under* the DEFINE SGROUP statement overrides any IMODE value specified *on* the DEFINE SGROUP statement. Also, an IMODE parameter in a DEFINE SGROUP section overrides all IMODE values in the DEFINE CGROUP section and in the DEFINE PARMS section.

```
IMODE=mode
```

Valid values:

DM	Indicates that CA-ASTEX is to run in *detail mode* to continuously collect performance data for the storage group.
EM	Indicates that CA-ASTEX is to run in *exception mode* to collect performance data for all members of the storage group only if the response time objective is not being met. Exception mode is useful only when you specify it for *all* volumes that share the same I/O path. Specifying exception mode for individual volumes will not help you find and solve path-related problems.
NONE	Indicates that CA-ASTEX should not collect any performance data for the storage group.

Example: Since volume SYSRR1 already has an IMODE value assigned to it in the 3390S contention group, you should not define a different value for it in the "SYSTEM" storage group:

```
DEFINE CGROUP,NAME=3390S
       VOL=SYSRR1,IMODE=DM
*
*
DEFINE SGROUP,NAME=SYSTEM
       VOL=(SYS*)
```

For more detailed information about the IMODE parameter, see the sections "DYNUNIT and the SMS DATACLAS Constructs" and "Making Entries in the DEFINE SGROUP Section" of this chapter.

MAYMSR Parameter

This parameter defines your I/O millisecond response time objective for data sets defined to the SMS may-cache storage class type.

Note: A MAYMSR parameter specified *below* the DEFINE SGROUP statement overrides any MAYMSR value specified *on* the DEFINE SGROUP statement. Also, any MAYMSR parameter in a DEFINE SGROUP section overrides all MAYMSR values in the DEFINE CGROUP and DEFINE PARMS sections.

```
MAYMSR=ttt
```

Valid values:

ttt	Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.

Default: `MAYMSR=25`

MSROBJ% Parameter

This parameter defines the percentage of time your response time objective should be met.

Note: A MSROBJ% parameter specified *below* the DEFINE SGROUP statement overrides any MSROBJ% value specified *on* the DEFINE SGROUP statement. Also, any MSROBJ% parameter in a DEFINE SGROUP section overrides all MSROBJ% values in the DEFINE CGROUP and DEFINE PARMS sections.

```
MSROBJ%=pp
```

Valid values:

pp	Identifies a percentage value between 50 and 99%.
Default:	`MSROBJ%=90` This default millisecond response time objective indicates that 90% of the I/O response times should be less than or equal to the values you specify on the NSMSMSR, MUSTMSR, MAYMSR, and NEVMSR parameters.

MUSTMSR Parameter

This parameter defines your I/O millisecond response time objective for data sets defined to the SMS must-cache storage class type.

Note: A MUSTMSR parameter specified *below* the DEFINE SGROUP statement overrides any MUSTMSR value specified *on* the DEFINE SGROUP statement. Also, any MUSTMSR parameter in a DEFINE SGROUP section overrides all MUSTMSR values in the DEFINE CGROUP and DEFINE PARMS sections.

```
MUSTMSR=ttt
```

Valid values:

ttt	Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.
Default:	`MUSTMSR=10`

NEVMSR Parameter

This parameter defines your I/O millisecond response time objective for data sets defined to the SMS never-cache storage class type.

Note: A NEVMSR parameter specified *below* the DEFINE SGROUP statement overrides any NEVMSR value specified *on* the DEFINE SGROUP statement. Also, any NEVMSR parameter in a DEFINE SGROUP section overrides all NEVMSR values in the DEFINE CGROUP and DEFINE PARMS sections.

```
NEVMSR=ttt
```

Valid values:

ttt Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.

Default: `NEVMSR=35`

NSMSMSR Parameter

This parameter defines your I/O response time objective for volumes not defined to SMS.

Note: A NSMSMSR parameter specified *below* the DEFINE SGROUP statement overrides any NSMSMSR value specified *on* the DEFINE SGROUP statement. Also, any NSMSMSR parameter in a DEFINE SGROUP section overrides all NSMSMSR values in the DEFINE CGROUP and DEFINE PARMS sections.

```
NSMSMSR=ttt
```

Valid values:

ttt Identifies a maximum time limit between 0 and 999 milliseconds. You want I/O response times to be *less than or equal to* this value.

Default: `NSMSMSR=30`

This default response time objective indicates that 90% of the I/O response times should be less than or equal to 30 milliseconds each. (The MSROBJ% parameter would also have to be set to `90`.)

VOL Parameter

This parameter lists members of a storage group by their *volume serial numbers* (volsers). You can repeat the VOL parameter as many times as necessary and define any number of volumes on a single VOL parameter.

Note: You *cannot* specify both VOL and DUA parameters (see above) under a single DEFINE SGROUP statement, use one or the other.

```
VOL=(vvvvvv)
```

Valid values:

vvvvvv Specifies a volume serial number. If you want to define several volumes with similar volsers, you can substitute an asterisk in place of specific characters.

Do not assign a specific volser (or corresponding DUA) to more than one storage group. CA-ASTEX will not analyze storage groups that have members in common.

Example:

```
DEFINE SGROUP,NAME=IMS
       VOL=(VOL*)
```

For the Migration Manager, volumes specified in a DEFINE SGROUP that are also defined in HSM ADDVOL commands in the ARCCMD*xx* member of the HSM PARMLIB will be added to the storage group named in the DEFINE SGROUP. Warning messages are issued to the JOURNAL DD, indicating the volumes were not added to the <ML1>, <NONSMS>, or <NONHSM> storage group. This allows you to group volumes as desired.

CA-ASTEX Server Parameters

The CA-ASTEX Server acts as an SNMP agent to send configuration and performance information to BrightStor DASD Manager for Unicenter and CA-ASTEX (formerly Unicenter TNG DASD Manager Option). This section tells you how to set up the CA-ASTEX Server address space parameters. These parameters are separate from the basic PMC parameters. They can be defined in either a sequential data set, or in a member of a partitioned data set.

Rules

Regardless of the data set you choose, you must specify your parameters according to these rules:

- The syntax for the server parameters is

  ```
  keyword=value
  ```

  Example: `COMPLEXID=SOMEPLEX`

- A * in column 1 signifies a comment.

  Example: `* this is a comment line`

- Comments may be placed after a parameter declaration, separated from the declaration by a blank.

  Example: `DBID=123 A comment may be placed here`

The Parameters

The following parameters apply to the Server component of CA-ASTEX.

Note: You **must** specify a value for parameter SNMP_ADDR=. We suggest that you review the values and defaults for all parameters.

COMPLEXID	PROBLEM_PERCENT
DBID	SNMP_ADDR
DINTVL	

The CA-ASTEX Server parameters are described below, listed in alphabetical order.

COMPLEXID Parameter

This parameter is the same as the basic PMC `COMPLEXID` parameter. It provides the CA-ASTEX Server with a logical name that it can use to group those systems that are part of a complex. All systems that are part of the same complex must use the same COMPLEXID.

```
COMPLEXID=xxxxxxxx
```

Valid values:

xxxxxxxx	A name consisting of at least one, but no more than 8 alphanumeric characters. The name must begin with an alpha character.
Default:	`COMPLEXID=_____` (eight underscore characters)

DBID Parameter

This parameter specifies the CA-Datacom database ID that the server will use. If you are running CA-Datacom/AD, the default of 620 should be sufficient because it is reserved for CA-ASTEX. If you are running CA-Datacom/DB, and you placed the CA-ASTEX database under a multi-user facility shared by other non-CA products, you may need to specify the database ID under which you installed CA-ASTEX.

```
DBID=xxx
```

Valid values:

xxx	21 to 999
Default:	`DBID=620`

DINTVL Parameter

This parameter corresponds to the DINTVL parameter of the main PMC parameters. Since all agents in a complex must run at the same interval, this value is the same value as the DINTVL parameter set for all CA-ASTEX agents in the complex.

```
DINTVL DINTVL=tttt
```

Valid values:

tttt	10 to 1440
Default:	`DINTVL=60`

PROBLEM_PERCENT Parameter

This parameter controls the number of problems found and the number of alerts generated. The higher the value of this parameter, the greater the percentage of problem I/Os found before an alert is generated for an LCU or volume. For example, setting PROBLEM_PERCENT to 10 will generate many alerts (lower threshold), and setting it to 65 will find very few alerts.

```
PROBLEM_PERCENT=pppPROBLEM_PERCENT
```

Valid values:

ppp	0 to 100
Default:	`PROBLEM_PERCENT=25`

SNMP_ADDR Parameter

This parameter specifies the host name or Internet address of the system where the IBM SNMP agent is running.

```
SNMP_ADDR=snmp address
```

Valid values:

snmp address	The host name or internet address of the system where the IBM SNMP agent is running.
Default:	None

Appendix

A

PMC: Calculating Storage and Space Requirements

CA-ASTEX requires common storage for its interval data collection areas and DASD space to store interval and trend data for PMC. A CLIST that calculates these storage and space requirements is located in member ASXCALC of the ASX.CNTL data set. If you prefer to do the calculations yourself, use the formulas on the following pages.

Calculating Your Common Storage Needs

The amount of common storage space that is necessary depends upon the number of volumes you want CA-ASTEX to monitor.

Use the formula in the following table to calculate your approximate common storage needs:

Space required for:	Number of bytes:
1. CA-ASTEX exit routines and communication areas	17,500
2. CA-ASTEX Allocation Director exits	2,200
3. The CA-ASTEX volume table Allow 352 bytes for each volume you want monitored. (Number of volumes) X 352=	____
4. Storage and contention group tables Allow 160 bytes for each storage group and 104 bytes for each contention group you define (or is defined by CA-ASTEX). (Number of storage groups) X 160 = (Number of contention groups) X 104=	 ____ ____

Space required for:	Number of bytes:
5. The volume table index Allow 1,024 bytes plus an additional 1,024 for each set of DASD devices that have a different value in the first two bytes of the DUA.	____
6. CA-ASTEX control blocks (188,000 bytes) with cache optimization. (8,000 bytes) **Note:** Cache optimization can use up to 100 KB of page-fixed extended CSA. Initial estimates are for an average of 20 KB.	____
Total: (Sum of all the above)	____

Important! *For items 3 and 5, CA-ASTEX uses extended common storage* (above the 16 megabyte line) *if it is available. For exit routines, communication areas (item 1), and group tables (item 4), CA-ASTEX does not use extended storage.*

How CA-ASTEX Reuses Common Storage

By specifying REUSE=YES in your start-up procedure, you indicate that CA-ASTEX should reuse as much common storage as possible each time you stop and restart PMC.

Every time you restart CA-ASTEX, the following common storage is reused:

- Volume table
- Storage for the Cache Optimizer

CA-ASTEX does ***not*** reuse the following common storage:

- Communication area storage
- Storage and contention group tables
- Volume table index storage

Calculating Space for the CA-ASTEX Interval Data

CA-ASTEX collects performance data for selected DASD volumes and optionally about the most active data sets and jobs on each volume. It collects this data for a certain interval of time (the number of minutes specified on the DINTVL parameter). Then, at the end of that interval, it optionally writes that data to an Interval Database, your SMF data set, or both (as specified on the DSNSMFT, JOBSMFT, IDB, and VOLSMFT parameters).

When writing performance data to one of these destinations, CA-ASTEX uses three different *record types*:

Data record type:	For performance information about:
Volume record	Volume I/O
Data set record	Data set I/O
Job record	Job I/O

Estimating Record Lengths

For each record type, you can estimate its length in bytes or, in other words, how much space it takes up in an SMF data set or Interval Database.

As a general rule, you need two 3390 cylinders per 150 volumes for each interval. (This assumes an average of 20 active data sets per volume and 10 active jobs.) You must allocate enough space for at least four intervals.

Note: The maximum size of the Interval Database should not exceed 2910 cylinders.

Length of the Volume Record

The following table shows how to calculate the length of the volume record:

What to calculate:	Number of bytes:
1. Length of the record header	84 bytes
2. Number of volumes Allow 350 bytes for each volume. (Number of volumes) X 350 =	____ bytes
Total length: (Sum of 1 and 2 above) = Total volume bytes	_____ total
Note: If you are recording to SMF and the total length of the volume record is greater than 32756, CA-ASTEX writes two or more data records.	

Length of the Data Set Record(s)

The following table shows how to calculate the length of the data set record(s):

What to calculate:	Number of bytes:
1. Length of the record header	84 bytes
2. Average number of active data sets per volume[1] Multiply this average by 350. (Average active data sets) X 350 =	____ bytes
3. Average length of one data set record (Sum of 1 and 2 above)	____ bytes
Total length of all data set records: Multiply the average data set record length (from line 3 above) by the number of volumes being monitored (Avg. length) X (no. of volumes being monitored)=Total data set bytes	_____ bytes

[1] To estimate this, look at the "max" value on the MLSIZE parameter to find the maximum number of data sets for each volume. Later, after you have started CA-ASTEX, you can look at a Problem Analysis Screen that has a summary/detail combination of "Volume" / "Data Set," estimate the average number of active data sets per volume, then reallocate the space, if necessary.

Length of the Job Record(s)

The following table shows how to calculate the length of the job record(s):

What to calculate:	Number of bytes:
1. Length of the record header	84 bytes
2. Average number of active jobs per volume[2] Multiply this average by 350. (Average active jobs) X 350 =	____ bytes
3. Average length of one job record (Sum of 1 and 2 above)	____ bytes
Total length of all job records: Multiply the average job record length (from line 3 above) by the number of job records written at the end of the data collection interval. (Avg. length) X (no. of volumes being monitored) = Total job bytes	_____ bytes

Total Cylinders Required for the Interval Database

The following table shows how to calculate the total number of cylinders required for the Interval Database:

What to calculate:	Result:
1. Total length of volume, data set, and job records (total vol bytes) + (total dsn bytes) + (total job bytes)=	_____ bytes
2. Total number of intervals* (number of days) X (number of intervals per day) =	_____ intervals
3. Bytes required to support the number of intervals you want in the Interval Database (number of intervals) X (total bytes) =	_____ bytes

---

[2] To estimate this, look at the "max" value on the MLSIZE parameter to find the maximum number of jobs for each volume. Later, after you have started CA-ASTEX, look at a Problem Analysis Screen that has a summary/detail combination of "Volume" / "Job," estimate the average number of active jobs on each volume, and reallocate the space, if necessary.

4. Number of cylinders required	
5. Divide total bytes for all intervals (item 3 above) by the number of bytes per cylinder for the device type the Interval Database resides on. This is a function of block size, and the following divisors can be used: ■ For 3380 models, use 641400 as the divisor. ■ For 3390 models, use 744480 as the divisor. (Byte value from item 3) / bytes per cylinder = *An Interval Database (IDB) must contain at least four intervals.	_____ cylinders

Note: Contiguous storage/space is not needed for the Interval Database.

Calculating Space for the Trend Database

The Trend Database is built from data in the Interval Database or the SMF data set. The Storage Performance Expert, Allocation Director, and Cache Management Expert all use trend data.

The following table shows how to find the amount of space to allocate:

What to Calculate:	Result:
1. Let "*x*" equal the maximum number of data sets per volume for which CA-ASTEX should maintain data in the Trend Databse. Multiply "*x*" by 320, and add 288 to the result. 320(*x*) + 288 = bytes	Total A: _____
2. Let "*nnn*" equal the maximum number of volumes CA-ASTEX is monitoring. Add "1" to *nnn*, and multiply the result by "Total A." Total A(*nnn* + 1) = bytes	Total B: _____
3. Let "*y*" equal the number of intervals per day (one 24-hour interval, two 12-hour intervals, three 8-hour intervals, and so on) that you want recommended solutions for. Add "1" to *y*, and multiply by "Total B" to get the number of bytes of space you need to allocate.	Total C:

	Total B(y + 1) = bytes	_____
4.	Let "f" equal the bytes per cylinder for your device type. Divide total bytes by bytes per cylinder and add "1" to the result. Total C / f + 1 = cylinders "Total D" is the number of cylinders you need to allocate.	Total D: _____

Appendix

B

MMC: Calculating Storage and Space Requirements

CA-ASTEX requires common storage to store data for MMC. A CLIST that calculates these storage and space requirements is located in member ASXCALC of the ASX.CNTL data set. If you prefer to do the calculations yourself, you find detailed formulas on the following pages.

Reviewing Your Common Storage Needs

The amount of required common storage space is fixed.

The following table shows your approximate common storage needs:

Common Storage Requirements	
Space required for:	**Number of bytes:**
MMC exit routines and communication areas	28672
Total:	28672

Note: The Migration Manager reuses all its common storage.

Calculating Space for the MMC Database

The Migration Manager Database comprises five data sets. The following sections describe each data set and how to calculate the size of the data set records.

Size of the prefix.DBMAIN Data Set

The *prefix*.DBMAIN data set contains information about each storage group and volume that the Migration Manager monitors.

The database technology used in all of CA-ASTEX "wraps" when it runs out of space. Much like a GDG with the LIMIT parameter, the "oldest" data is discarded to make room for the newest data. The difference here is that the space utilization takes place within an extent of space controlled internally by CA-ASTEX (to OS/390, this is a BDAM data set).

prefix.DBMAIN contains all the summary information, meaning there is:

- 1 record per storage group
- 1 record per volume
- 1 record for the entire DASD farm (the complex).

To allow for the long-term retention of this summary data without wrapping, we recommend you choose a large allocation quantity so you can track how MMC and HSM statistics have improved over the months. For example, to store a year of daily summary statistics, specify **365** for the number of cycles.

The following table shows how to calculate the approximate number of bytes and cylinders needed for this data set:

What to calculate:	**Result:**
1. Length of the complex record	713 bytes
2. Number of storage groups to be monitored. Multiply this number by 401. Groups X 401 =	____ bytes
3. Number of volumes to be monitored. Multiply this number by 399. Volumes X 399 =	____ bytes
4. Sum of 1, 2, and 3 above	____ bytes
5. **Total length of all data set records:** Multiply the size of one cycle (from line 4 above) by the number of cycles. (Cycles represents the number of cycles of information to keep. We recommend a minimum of 30.) (Cycle size) X (no. of cycles) = Total bytes	____ bytes
6. Number of cylinders required. Divide the total bytes for all records (item 5 above) by the number of bytes per cylinder for the device type on which the MMC Database resides. This is a function of block size, and the following divisors can be used: – For 3380 models, use 641400 as the divisor. – For 3390 models, use 744480 as the divisor. (Byte value from item 5) / bytes per cylinder =	_____ cylinders

Size of the prefix.DBFILE Data Set

The *prefix*.DBFILE data set contains information about each data set that the Migration Manager monitors. This includes data sets on primary DASD (ML0) volumes and compressed DASD (ML1) volumes. It also includes data sets on compressed tape (ML2) volumes, if users recall them.

Unlike *prefix*.DBMAIN, which can be calculated quite precisely using worksheets and/or the CLIST, space for *prefix*.DBFILE can only be estimated, because it is impossible to know in advance the amount of HSM activity and the number of data sets for which MMC makes recommendations.

prefix.DBFILE contains detail data (one record for every data set on ML0, ML1, and some data sets on ML2). Therefore, it needs to be quite large on a relative basis for the *current interval*. If you can estimate the number of data sets on these locations, the worksheet or CLIST is quite useful.

However, when a current interval is replaced by data from a new inventory of the DASD complex, it is aged and reduced in size. Data sets having any HSM activity are kept; data sets recommended by MMC for migration—or recall—are also kept. Because all other data sets in a current interval are dropped to trim the amount of new space used in *prefix*.DBFILE, the rate of growth in this data set cannot be accurately predicted.

Therefore, when the worksheet and/or CLIST prompts you for the number of cycles for *prefix*.DBFILE, you should understand that this answer *underestimates* the previous cycles of detail data you are able to keep. That is, if you specify 5 cycles, you may actually be able to store, for example, 15 cycles of detail data. Again, this is because the current interval is large and a function of the number of data sets, while the previous intervals are much smaller and a function of the amount of HSM activity and MMC recommendations.

We recommend you *keep the number of cycles relatively small*. You do not need detail data kept for many months, as we recommend for *prefix*.DBMAIN. The issue is how far back in time you might want to go in researching a particular HSM (migrate or recall) activity or for reporting purposes. Typically, we believe this would be only a few days; hence, we recommend specifying 30 cycles.

The following table shows how to calculate the approximate number of bytes and cylinders needed for this data set:

What to calculate:	Result:
1. Number of data sets to be monitored. Multiply this number by 219. Data sets X 219 =	____ bytes
2. Number of data sets to be monitored. Multiply this number by 25.	

What to calculate:	Result:
Data sets X 25 =	____ bytes
3. Number of cycles. (Cycles represents the number of cycles of information to keep. We recommend a minimum of 30.) Multiply this number by 2 above. (no. of cycles) X (2 above) =	____ bytes
4. **Total length:** (Sum of 1 and 3 above) = Total bytes	____ bytes
5. Number of cylinders required. Divide the total bytes for all data set records (item 4 above) by the number of bytes per cylinder for the device type on which the MMC Database resides. This is a function of block size, and the following divisors can be used: – For 3380 models, use 641400 as the divisor. – For 3390 models, use 744480 as the divisor.	_____
(Byte value from item 4) / bytes per cylinder =	cylinders

Size of the prefix.DBHIST Data Set

The *prefix*.DBHIST data set contains historical data set access patterns based on both the data set name and the access patterns within each storage group.

The following table shows how to calculate the approximate number of bytes and cylinders needed for this data set:

What to calculate:	Result:
1. Number of unique high-level data set nodes (or, prefixes) defined. Multiply this number by 98. Nodes X 98 =	____ bytes
2. Number of storage groups to be monitored. Multiply this number by 3677. Groups X 1480 =	____ bytes
3. **Total length:** (Sum of 1 and 2 above) = Total bytes	____ bytes
4. Number of cylinders required. Divide the total bytes for all data set records (item 3 above) by the number of bytes per cylinder for the device type on which the MMC Database resides. This is a function of block size, and the following divisors can be used: – For 3380 models, use 641400 as the divisor.	

– For 3390 models, use 744480 as the divisor.	_____
(Byte value from item 3) / bytes per cylinder =	cylinders

Size of the prefix.DBMISC Data Set

The *prefix*.DBMISC data set contains information about each management class defined. It also serves as a temporary holding area for data sets that are no longer defined on the complex being monitored.

The following table shows how to calculate the approximate number of bytes and cylinders needed for this data set:

What to calculate:	Result:
1. Number of management classes defined. Multiply this number by 21. Classes X 21 =	____ bytes
2. Number of data sets to be monitored. Multiply this number by 44. Data sets X 44 =	____ bytes
3. **Total length:** (Sum of 1 and 2 above) = Total bytes	____ bytes
4. Number of cylinders required. Divide the total bytes for all records (item 3 above) by the number of bytes per cylinder for the device type on which the MMC Database resides. This is a function of block size, and the following divisors can be used: – For 3380 models, use 641400 as the divisor. – For 3390 models, use 744480 as the divisor.	_____
(Byte value from item 3) / bytes per cylinder =	cylinders

Size of the prefix.DBMOVE Data Set

The *prefix*.DBMOVE data set contains requests passed to HSM to move data sets based on Migration Manager recommendations. This data set is not used until you activate the Migration Manager for one or more storage groups.

The following table shows how to calculate the approximate number of bytes and cylinders needed for this data set:

What to calculate:	Result:
1. Maximum number of move requests passed to HSM each day. Multiply this number by 76. Maximum requests X 76 =	____ bytes
2. Number of cylinders required. Divide the total bytes for all records (item 1 above) by the number of bytes per cylinder for the device type on which the MMC Database resides. This is a function of block size, and the following divisors can be used: – For 3380 models, use 641400 as the divisor. – For 3390 models, use 744480 as the divisor.	_____
(Byte value from item 1) / bytes per cylinder =	cylinders

Appendix

C

PMC Technical Descriptions

Initialization Functions

CA-ASTEX performs the following initialization functions for PMC when you issue the S ASXPMC command:

1. It reads the CA-ASTEX parameter data set. This data set defines which volumes CA-ASTEX should monitor. It also identifies the *contention groups, storage groups*, *ESCON groups,* and *logical volume groups* these volumes belong to.

   - Contention groups contain volumes that share common paths. By default, CA-ASTEX assigns volumes to contention groups in both single- and multiple-CPU environments. The contention group names are based on the Storage Director IDs in the Storage Director of the associated DASD controller. In multiple-CPU environments, the contention groups are assigned across all CPUs.
   - Storage groups contain volumes that have common functions. If you are using SMS, CA-ASTEX automatically assigns volumes to the correct SMS storage groups. If you are not using SMS, you must define your own storage groups in the DEFINE SGROUP section of your parameter data set.
   - An ESCON group contains volumes on channel paths that share a common ESCON director. CA-ASTEX does not automatically assign volumes to ESCON groups. You must define your own ESCON groups in the DEFINE ESCON section of your parameter data set. You do not define ESCON groups if you have ESCON channels but no ESCON directors.
   - RAID storage architectures use various mapping strategies to define logical volumes to physical disks. Logical volumes can be defined to a single physical disk or multiple physical disks. Multiple logical volumes can also be defined to a single physical disk. CA-ASTEX automatically determines which logical volumes are mapped or defined to the same physical disk(s) and forms a unique name for each group of the logical volumes. This function is available for only EMC Symmetrix ICDAs and IBM RAMAC Array Subsystems.

2. It dynamically installs CA-ASTEX I/O exit routines, Allocation Director exit routines, and the configuration change exit routines, and moves them into common storage. (Configuration change exit routines are exits that dynamically detect when volumes have been added or deleted—brought online or offline— to the system. These exits are removed when you stop CA-ASTEX.)

   The CA-ASTEX I/O and Allocation Director exit routines remain in place after CA-ASTEX is stopped. To remove them, you must IPL the system.

3. It obtains additional common storage areas for the following:
   a) A volume table where CA-ASTEX maintains I/O information for each volume monitored
   b) Group tables where it maintains measurements of storage, contention, and logical volume group I/O activity
   c) A communication area where CA-ASTEX maintains status indicators, parameter values, and so on
   d) Cache optimization control information

I/O Response Time Measurements

CA-ASTEX calculates I/O response time by adding several elapsed time values:

Queue time	How long an I/O request waits in the I/O subsystem queue for a DASD that is busy with another I/O operation
Pending time	How long an I/O request waits in a channel subsystem queue for a DASD or path that is busy with another I/O operation. In traditional DASD environments, pending time accumulates until the requested physical device is available. In RAID environments, pending time accumulates until the logical device is available. The end of pending time accumulation no longer guarantees that the physical device is available in a RAID environment.
Connect time	How long the DASD is “busy” searching for or transferring the requested data
Disconnect time	How long the DASD is “busy” but not actually searching for or transferring requested data. With RAID storage environments, disconnect time can now include a portion of traditional pending time. This occurs when pend time accumulation stops because the logical volume is available, but the I/O is unable to continue because the physical device(s) is not available. The time an I/O waits for the physical device(s) to become accessible is accumulated as disconnect time.
Dispatch time	How long the I/O request waits for IOS, channel subsystem dispatching, or I/O interrupt processing

Storage Classes

This section shows how CA-ASTEX determines an I/O's storage class type on 3990-type and 3880-type controllers.

3990-type Controller

The following table shows how CA-ASTEX classifies I/Os on a 3990-type controller. The table lists SMS settings at I/O time and how CA-ASTEX then classifies that type of I/O.

Is volume turned on to cache?	Is CASF token valid?	What is the CASF token?	I/O type?	Read Bias (RB) or Write Bias (WB)?	Direct or Sequential?	DASD Fast Write?	CA-ASTEX assigns this storage class.
Yes	Yes	Must	Read	RB or no RB	Direct or Sequential	No	Must
Yes	Yes	May	Read	RB or no RB	Direct or Sequential	No	May
Yes	Yes	Never	Read	RB or no RB	Direct or Sequential	No	Never
No	Yes	Must	Read	RB or no RB	Direct or Sequential	No	Must
No	Yes	May	Read	RB or no RB	Direct or Sequential	No	May
No	Yes	Never	Read	RB or no RB	Direct or Sequential	No	Never
Yes	Yes	Must	Write	WB or no WB	Direct or Sequential	No	Must
Yes	Yes	May	Write	WB or no WB	Direct or Sequential	No	May
Yes	Yes	Never	Write	WB or no WB	Direct or Sequential	No	Never
Yes	Yes	Must	Write	WB	Direct	Yes	Must
Yes	Yes	Must	Write	no WB	Direct	Yes	May
Yes	Yes	May	Write	WB or no WB	Direct	Yes	May
Yes	Yes	Never	Write	WB or no WB	Direct	Yes	Never

Is volume turned on to cache?	Is CASF token valid?	What is the CASF token?	I/O type?	Read Bias (RB) or Write Bias (WB)?	Direct or Sequential?	DASD Fast Write?	CA-ASTEX assigns this storage class.
No	Yes	Must	Write	WB or no WB	Direct or Sequential	No	Must
No	Yes	May	Write	WB or no WB	Direct or Sequential	No	May
No	Yes	Never	Write	WB or no WB	Direct or Sequential	No	Never
No	Yes	Must	Write	WB	Direct	Yes	Must
No	Yes	Must	Write	no WB	Direct	Yes	May
No	Yes	May	Write	WB or no WB	Direct	Yes	May
No	Yes	Never	Write	WB or no WB	Direct	Yes	Never
Yes	No	----	Read	----	----	No	Must
Yes	No	----	Write	----	----	Yes	May
Yes	No	----	Write	----	----	No	Never
No	-	----	R/W	----	Direct or Sequential	-	Never
No	No	----	R/W	----	Direct or Sequential	-	Never

3880-type Controller

The following table shows how CA-ASTEX classifies I/Os on a 3880-type controller. The table lists SMS settings at I/O time and how CA-ASTEX then classifies that type of I/O.

Is controller turned on to cache?	What is the CASF token?	I/O type?	CA-ASTEX assigns this storage class.
Yes	----	Read/Write	Must
Yes	----	Read/Write	Must
No	----	Read/Write	Never

I/O Exit Processing

This is a summary of the processing done by the CA-ASTEX I/O exit routines:

1. Each I/O request is analyzed to see if it qualifies for CA-ASTEX processing. The I/O must be issued against a volume that CA-ASTEX is monitoring.
2. If cache optimization is active for the contention group, CA-ASTEX determines whether the track is a good cache candidate. If the track is not a good cache candidate, CA-ASTEX inhibits cache access.
3. For I/Os that qualify, CA-ASTEX measures the following times:
   a) The time the user program issued the I/O request (start of queue time)
   b) The time the I/O actually started on the device (end of queue time). This value reflects the time when the DASD device accepts the first channel command.
   c) The I/O completion time
4. At I/O completion, the I/O response time is calculated, and the CA-ASTEX volume and group tables are updated.
5. If the I/O response time exceeds the user-defined response objective, additional statistics are updated in the volume and group tables to indicate the cause of the problem.
6. The data set name, job name, program name and additional data for the I/O are passed to the CA-ASTEX address space for processing in the following cases:
   a) If the originating group or volume is being traced
   b) If the originating group or volume is set to detail mode
   c) If the originating group is set to exception mode and the percent of unacceptable I/Os for the group is exceeded

Address Space Processing

This is a summary of the processing done by the CA-ASTEX performance address space code:

1. A STIMER loop is established in order to process I/O trace data passed by the I/O exit routines. At each "timer pop," CA-ASTEX processes I/O data that has been passed to the CA-ASTEX address space routines. This processing includes:
   a) Writing trace records for volumes that are being traced
   b) Updating the managed lists for volumes or groups that are set to detail mode
   c) Writing records to an SMF data set and/or the Interval Database at the end of the current CA-ASTEX interval
2. A VTOC reader subtask is attached. For certain system data sets (LINKLIST, catalogs, and so forth) the data set name is unavailable to the I/O exit routines. The VTOC reader subtask reads the VTOC and maintains a list of the most active VTOC entries.
3. A dynamic change subtask is attached. This subtask dynamically detects any changes to your configuration including cache status, addition or deletion of volumes, and storage group changes.
4. Optionally, an alert subtask is attached to provide configuration and alert information to the CA-ASTEX server. This information is ultimately displayed via the Unicenter TNG DASD Manager Option.
5. CA-ASTEX commands are processed.

Data Collection Functions

CA-ASTEX continuously collects performance data for groups (either contention or storage groups) and volumes. The group and volume data it collects reflects performance either for the entire day or else for a certain interval of time (from 9:00 a.m. to 10:00 a.m., for example).

CA-ASTEX also can collect detailed job and data set data that reflects performance just for a certain interval of time. However, CA-ASTEX collects this data only if it is operating in one of the following two **data collection modes**:

1. **Detail mode** - CA-ASTEX continuously gathers detailed data set and job data for particular volumes.
2. **Exception mode** - CA-ASTEX gathers data set and job data only when performance is poor (when I/O response times are too long).

How Exception Mode Works

When running in exception mode, CA-ASTEX automatically switches in and out of detail mode. As long as the I/O response times for a particular volume or group of volumes meet your objectives, CA-ASTEX stays out of detail mode.

But, when conflicting I/O requests cause poor response for data sets on a particular volume, CA-ASTEX switches into detail mode and collects detailed job or data set data for that volume. Likewise, if contention occurs among volumes that share a single I/O path (volumes in the same contention group), CA-ASTEX switches into detail mode and collects detailed job or data set data for all volumes in the group.

Thus, when CA-ASTEX runs in exception mode the amount of detailed data set or job data you receive depends upon how many periods of poor response it detects.

Note: Although you define a "default" data collection mode (on the IMODE parameter) when you install CA-ASTEX, you also can change the data collection mode any time you like by issuing the CA-ASTEX SGRP and SVOL commands. See the chapter "Operating CA-ASTEX" in the CA-ASTEX *User Guide* for more information on controlling the collection of data.

Managed List Usage

CA-ASTEX limits the amount of job and data set performance data that it collects and later screens by maintaining a series of **managed lists**. To maintain and display detailed job and data set data for every volume that CA-ASTEX monitors in detail mode would require too many resources. So, instead, CA-ASTEX saves only the most significant job or data set information from the current detail mode interval. (This interval is a length of time specified on the DINTVL parameter.) CA-ASTEX saves this information in a "managed list," a list that it continually updates by adding entries for the most active data sets and jobs or deleting entries for data sets and jobs that become inactive.

Each managed list never contains entries for less than the *min* or more than the *max* number of data sets or jobs specified on the MLSIZE parameter in the ASXPARM member of ASX.CNTL:

- If a volume has the *min* number of active data sets, the managed list contains entries for at least that *min* number of data sets.
- If a volume has more than the *max* number of active data sets, the managed list contains entries for only the *max* number of busiest ones.
- If a volume has more than the *min* number of data sets, but less than the *max*, the managed list contains entries for data sets (or jobs) that contribute to at least 90% of the volumes total I/O activity.

CA-ASTEX stops updating the managed lists when the detail mode interval (the DINTVL time) expires. Optionally, it writes the detailed job and data set data to an SMF data set and/or Interval Database where you save other CA-ASTEX data. Then, it purges each managed list and starts accumulating data for the next detail mode interval.

Page Data Set I/O Measurements

CA-ASTEX measures I/O requests for page data sets as follows:

- It updates its page data set I/O statistics when the paging channel program terminates.
- The statistics for page data sets reflect I/O activity for each SUSPEND/RESUME cycle. For example, if a channel program starts and is "suspended" and "resumed" three times before it terminates, then CA-ASTEX counts that as four I/Os.
- It does not update "Prob%" measurements for paging I/O.

VSAM Data Set Statistics

The CA-ASTEX online screens and batch reports contain the following types of performance information for VSAM data sets:

- Statistics for each VSAM component, on systems running DFP (Data Facility Product)
- Statistics for each VSAM cluster or path name, for systems not running DFP

For VSAM data sets defined with the imbed option, CA-ASTEX counts I/Os to the imbedded index *under the data set name of the index component*. You should consider this when reviewing the CA-ASTEX reports and recommendations.

ESCON Measurements and Analysis

CA-ASTEX provides comprehensive measurements in ESCON environments:

- At initialization, CA-ASTEX identifies the Logical Control Units that share the same ESCON directors by reading the DEFINE ESCON statement that you specify in the ASXPARM member of ASX.CNTL.
- When measuring each I/O, CA-ASTEX captures the ESCON director port busy associated with the I/O. A DPB of non-zero indicates that an I/O has been delayed in the ESCON Director. ESCON Director port busy occurs when the ESCON Director is busy and an I/O must be delayed until a port is free.

  If an RPS delay also occurs for the I/O, the RPS delay is attributed to the ESCON Director port busy condition. RPS delays can therefore be caused by ESCON Director port busy.
- When the Storage Performance Expert analyzes path problems, its analysis takes into account ESCON pathing delays in addition to RPS pathing delays. If ESCON delay problems have caused response time delays, the solutions provided by the Storage Performance Expert takes the ESCON directors into account. It recommends moving data sets or volumes to other paths (Logical Control Units) that do not share the same ESCON Director as the problem path. This resolves the problem by offloading the contention on the ESCON Director.

Cache Hit Identification

About 80 to 85% of I/O transactions use RPS channel programs. When data is read from cache memory during an I/O transaction, the RPS channel program finishes with a "disconnect time" of 0. When CA-ASTEX finds a disconnect time of 0, it treats the associated I/O transaction as a cache hit.

Sometimes, however, a disconnect time of 0 appears even if the I/O transaction was not a cache hit. Usually, this happens when a non-RPS channel program handles the I/O transaction. If the non-RPS channel program does not require a "seek," the resulting disconnect time may be zero, even if the I/O transaction was a cache miss rather than a cache hit. In such cases, CA-ASTEX has to use a value other than disconnect time to determine cache hits.

Note: A miss occurs when data is read from DASD rather than the cache buffer during an I/O transaction.

With non-RPS channel programs that do not require a seek, "connect time" is a better indicator of cache hits. Typically, a non-RPS channel program searches for a record (by performing, for example, a VTOC search, a catalog search, or a PDS directory search) when the device sector number is not known. If a cache hit occurs, connect time is usually very brief, for two main reasons. First, the search for that record (in the cache buffer) is conducted at the electronic speed of the cache microcode. Second, the amount of data transferred from the record is likely to be small. Therefore, CA-ASTEX identifies cache hits for non-RPS programs by looking for very low connect times.

Some of the newer RAID devices (RAID 1, RAID 5, and RAID 6) can have cache hits when disconnect time is greater than 0. CA-ASTEX takes this into account when determining cache hits for these controllers.

CA-ASTEX and Cache Controllers

CA-ASTEX is fully functional with any cache controller that supports the Extended Count Key Data (ECKD) architecture. Included in this architecture is the DEFINE EXTENT Channel Command Word (CCW). CA-ASTEX uses the information in the DEFINE EXTENT to determine an I/O's effect on a cache controller. CA-ASTEX also uses the SENSE SUBSYSTEM CCW to determine if a volume is turned on to cache.

Some cache controllers do not fully support ECKD and SENSE SUBSYSTEM CCW:

- Amdahl 6100 at microcode levels below 4.3
- Amdahl 6880
- IBM 9345
- EMC 4200 and 4400 series at microcode levels below 1709
- EMC 4800 series at microcode levels below 3709

If you have these controllers installed and you are using full volume caching, you must specify the CA-ASTEX NICACH parameter for each volume turned on to cache and NVS storage. All I/Os to the volumes you specify on the NICACH parameter are then categorized as cache candidates. This allows CA-ASTEX to produce cache measurements, displays, and expert analysis for those devices.

If you are using the AMDAHL SPMS facility to control caching at the data set level, keep the following in mind:

1. CA-ASTEX assumes that all data sets are using cache. Since the AMDAHL cache controllers determine a data set's access to cache resources, CA-ASTEX may wrongly assume that cache misses are causing cache loading. In turn, the CA-ASTEX Cache Statistics facility may incorrectly show high negative net benefit measurements and load counts for data sets not actually being cached.
2. The CA-ASTEX Cache Management Expert looks at negative net benefit to determine whether a cache user is good or bad. As a result, the Expert may identify data sets that are not actually using cache as poor cache users. The Expert then recommends that you move these data sets out of cache. A special APAR is available to eliminate most of these recommendations. Please contact Computer Associates Technical Support for more details.

Note: Keep in mind that the CA-ASTEX Cache Optimizer facility does not function on cache controllers that do not fully support ECKD architecture.

CA-ASTEX and RAID Devices

RAID (Redundant Array of Independent Disks) architectures offer improvements in performance and/or reliability. Typically, RAID configurations use cache and non-volatile storage to enhance performance and reliability and to mitigate the various penalties associated with RAID architectures.

CA-ASTEX version 2.8 supports RAID 1, RAID 5, and RAID 6 architectures. Brief explanations of each of these architectures follow:

1. RAID 1: This architecture supports a *one-to-one mapping* between the logical device, which is represented by an OS/390 UCB(unit control block), and the physical disk. When a write I/O is issued to a UCB located behind a RAID 1 controller, the I/O's data is written to the logical volume's corresponding physical disk (similar to 3380/3390 technology). Data redundancy is provided via "mirroring." EMC Symmetrix ICDAs implement a version of RAID 1.
2. Enhanced RAID 1: This architecture allows multiple logical volumes to be mapped to a single physical disk, creating a *many-to-one mapping*. The group of logical volumes associated with a single physical disk is called a Logical Volume Group. The data associated with a write I/O is still written to a single physical disk, but I/Os from multiple logical volumes may contend for the same physical disk. "Mirroring" is implemented when data redundancy is required. The Hyper-volume extension feature available on EMC Symmetrix ICDAs is an example of an enhanced RAID 1 implementation.

3. RAID 5: This architecture creates a *one-to-many mapping* between a single logical device and a group of physical disks. Multiple logical volumes are associated with a specific group of physical devices. The set of logical volumes associated with a particular group of physical devices is called a Logical Volume Group (LVG). When a write I/O is issued to a UCB located behind a RAID 5 controller, the data associated with the I/O is written across the physical disks that are a part of the logical volume's Logical Volume Group. A parity disk is used to provide data redundancy. The implementation of the IBM RAMAC Array Subsystem, the IBM RAMAC Scalable Array (formerly STK Kodiak), and the HDS 7700 are all based on RAID 5 architecture.
4. RAID 6: This architecture supports a *one-to-many mapping* between a logical volume and the physical disks. When a write I/O is issued to a UCB located behind a RAID 6 controller, the I/O's data is written across multiple physical disks. Another differentiating characteristic of this RAID architecture is that data is never updated in place. Each write I/O invalidates the current data location and writes the data to a new location. Parity disk(s) provide data redundancy for RAID 6 implementations. A RAID 6 implementation example is the IBM RAMAC Virtual Array (formerly STK Iceberg).

Not all traditional I/O measurements can be applied to RAID 1, RAID 5, and RAID 6 architectures. These architectures have redefined meanings for old measurements.

RAID 1 architecture allows multiple logical volumes (OS/390 UCBs) to be mapped to a single physical disk. This configuration permits I/Os from multiple OS/390 UCBs to contend for the same physical device. To notify you of this problem, CA-ASTEX provides a problem measurement called Physical Disk Contention (PDC). This measurement is updated when an I/O request fails to meet its response time objective due to contention at the physical disk level.

Both RAID 5 and RAID 6 architectures map a single logical volume to multiple physical disks. Because multiple logical volumes are mapped to the same physical disks, multiple logical volumes can simultaneously issue I/Os that must access the same group of physical disks. The problem measurement used to identify this contention is Physical Disk Contention (PDC).

When the Storage Performance Expert finds that response time problems are occurring because of Physical Disk Contention, it recommends that data sets be moved to a volume that is part of another Logical Volume Group or to a volume that is not part of any Logical Volume Group.

In addition, the Cache Simulator supports RAID 1, RAID 5, and RAID 6 cache configurations. This allows you to see the impact of proposed cache changes to your RAID devices before you make them.

Appendix

D

Contacting Technical Support

Overview

This chapter contains information about:

- Identifying and resolving problems
- Contacting Computer Associates Technical Support
- Receiving ongoing product releases and maintenance
- Requesting product enhancements

Diagnostic Procedures

Refer to the flowchart below for a summary of the procedures you should follow if you have a problem with a Computer Associates software product. Each of these procedures is detailed on the following pages.

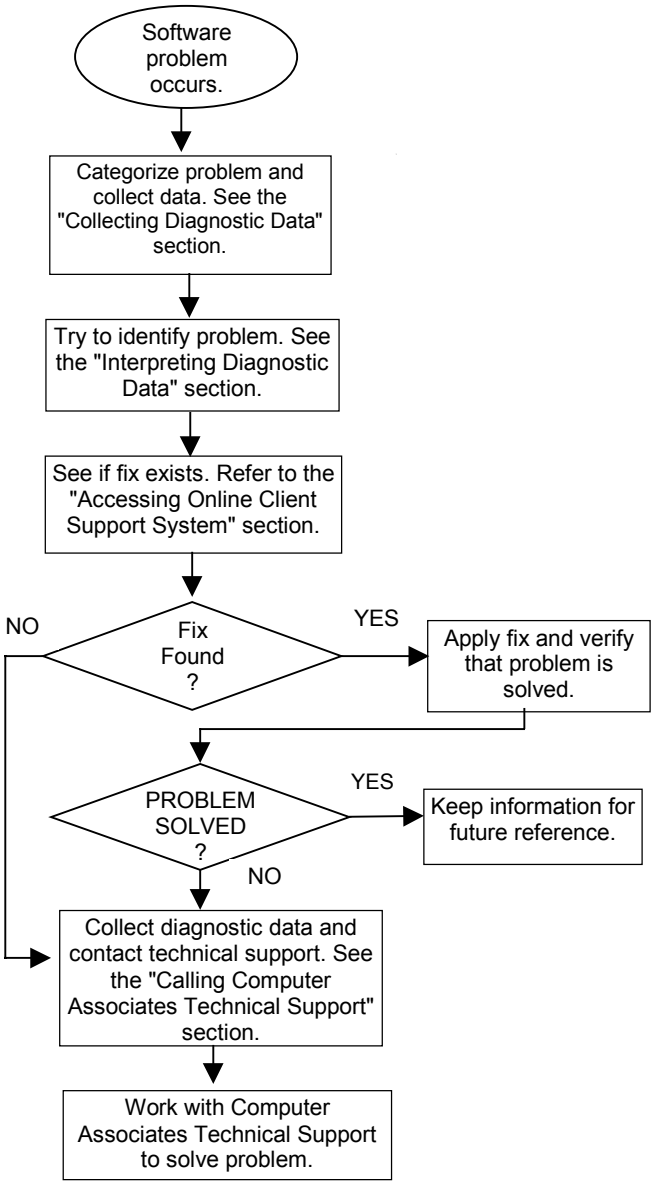

Collecting Diagnostic Data

The following information is helpful in diagnosing problems that might occur:

- Control statements used to activate your product
- JCL used to install or activate your product
- Relevant system log or console listings
- Relevant system dumps or product dumps
- List of other IBM or third-party products that might be involved
- Manufacturer, model number, and capacity of your hardware
- Numbers and text of IBM or CA error messages associated with the problem
- Names of panels where the problem occurs
- Listings of all fixes applied to all relevant software, including:
  - The date fixes were applied
  - Fix numbers
  - Names of components to which fixes were applied
- Short description of problems

Interpreting Diagnostic Data

When you have collected the specified diagnostic data, write down your answers to the following questions:

1. What was the sequence of events prior to the error condition?
2. What were the circumstances when the problem occurred and what action did you take?
3. Has this situation occurred before? What was different then?
4. Did the problem occur after a particular PTF was applied or after a new release of the software was installed?
5. Have you recently installed a new release of the operating system?
6. Has the hardware configuration (tape drives, disk drives, and so forth) changed?

From your response to these questions and the diagnostic data, try to identify the cause and resolve the problem.

Accessing the Online Client Support System

StarTCC, the web-based portion of CA-TCC (CA-Total Client Care), gives you real time, interactive access to Computer Associates product support information through the Internet. Using StarTCC, you can:

- Open new issues
- Browse or update your existing issues
- Perform keyword searches
- Download solutions, service packs, and important notices regarding Computer Associates products, maintenance, and documentation

Requirements for Using StarTCC

The following are the requirements to use StarTCC:

- You must be a CA client with a current maintenance agreement.
- You must register through the CA Internet site.
- You must access the Internet with a browser that supports the HTML specification 4.0 or higher, such as Netscape Navigator 4.0 or higher, or Microsoft Internet Explorer 4.0 or higher. (You can download one of these browsers from the Computer Associates Technical Support page at http://esupport.ca.com.)

  Browsers that meet the HTML requirement support the following functions, which are required for StarTCC:

  - Secure sockets layer (SSL) to encrypt your transaction traffic
  - Encrypted data records (known as COOKIES)
  - HTML tables

StarTCC Security

StarTCC runs as a secured server (SSL). You may need to configure your browser to enable SSL. Guidelines for doing this are provided on the CA Technical Support page.

Accessing StarTCC

To access StarTCC, access http://esupport.ca.com or click the Support button on the CA home page and follow the links for StarTCC.

If you are a first time user, you must register before you can access StarTCC online. Select the registration option to identify yourself to StarTCC. There are prompts for all required information, including your name, site ID, CA-StarTrak PIN, company name, e-mail address, postal address, and desired password for accessing StarTCC.

Note: If you do not have a CA-StarTrak PIN, StarTCC provides one for you when you register.

Select the access option to begin using StarTCC. When prompted, enter your user ID and password.

Accessing the Product Support Directory Online

The Computer Associates Product Support Directory lists each CA product and the telephone number to call for primary support for that product. To access the Product Support Directory online, click the Support button on the CA home page. Follow the links.

CA-TLC: Total License Care

Many CA software solutions use license keys or authorization codes to validate your hardware configuration. If you need assistance obtaining a license key or authorization code, contact the CA-TLC: Total License Care group at 1-800-338-6720.

Calling Computer Associates Technical Support

Computer Associates provides telephone support for all its products.

If you are in North America, refer to the *Product Support Directory* for the specific CA Technical Support phone number for each product. Outside North America, call your local Computer Associates Support Center during normal business hours.

Emergency phone numbers are also available for after-hours technical support. After hours calls should be limited to severity 1 problems.

Note: Only your local Computer Associates Support Center can provide native language assistance. Please use English when contacting any North American center.

If you are unable to resolve the problem, please have the following information ready before contacting Computer Associates Technical Support:

All the diagnostic information described in

- Collecting Diagnostic Data.
- Product name, release number, operating system and service pack.
- Product name and release number of any other software you suspect is involved.
- Release level of the operating system.
- Your name, telephone number, and extension (if any).
- Your company name.
- Your site ID.
- A severity code. This is a number (from 1 to 4) that you assign to the problem. Use the following to determine the severity of the problem:
    1. A "system down" or inoperative condition
    2. A suspected high-impact condition associated with the product
    3. A question concerning product performance or an intermittent low-impact condition associated with the product
    4. A question concerning general product utilization or implementation

Product Releases and Maintenance

Clients are requested to operate only under currently supported releases of the product.

Clients with current maintenance agreements also receive ongoing maintenance. When a new release of the system is available, a notice is sent to all current clients.

Requesting Enhancements

Computer Associates welcomes your suggestions for product enhancements. All suggestions are considered and acknowledged. You can use either of two methods to request enhancements:

- Access StarTCC from the CA web site to open an issue for the enhancement request. For details, see Accessing the Online Client Support System in this chapter.
- Contact your account manager or a technical support representative who will initiate a Demand Analysis Request (DAR) for you.

The DAR system serves as a central receiving point for all enhancement requests.

Index

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W